

Market Microstructure Invariance: Theory and Empirical Tests*

Albert S. Kyle
Robert H. Smith School of Business
University of Maryland
akyle@rhsmith.umd.edu

Anna A. Obizhaeva
Robert H. Smith School of Business
University of Maryland
obizhaeva@rhsmith.umd.edu

June 7, 2013

Abstract

Using the intuition that financial markets transfer risks in business time, we define “market microstructure invariance” as the hypothesis that the distribution of risk transfers (“bets”), transactions costs, resilience, and market efficiency are constant across assets when measured in units of business time. In calendar time, the invariance hypothesis results in specific empirically testable invariance relationships among those variables. A meta-model implies that invariance relationships are ultimately related to granularity of information flow. Based on a dataset of 400,000+ portfolio transition orders, we show that quantitative predictions of microstructure invariance concerning bets sizes and transactions costs as functions of observable volume and volatility closely match the data. We calibrate invariant parameters and discuss implications for financial markets.

Keywords: market microstructure, order size, number of orders, transactions costs, liquidity, resilience, efficiency, portfolio transitions, invariance, bid-ask spread, market impact.

*We are grateful to Elena Asparouhova, Peter Bossaerts, Xavier Gabaix, Lawrence Glosten, Pankaj Jain, Mark Loewenstein, Georgios Skoulakis, and Vish Viswanathan for helpful comments. Obizhaeva is also grateful to the Paul Woolley Center at London School of Economics for its hospitality as well as Ross McLellan, Simon Myrgren, Sébastien Page, and especially Mark Kritzman for their help.

Introduction

This paper¹ proposes a modeling principle for financial markets that we call “market microstructure invariance.” When portfolio managers trade financial assets, they can be modeled as playing trading games in which risks are transferred. Market microstructure invariance begins with the intuition that these risk transfers, which we call “bets,” take place in business time. The rate at which business time passes—market “velocity”—is the rate at which new bets arrive into the market. For actively traded assets, business time passes quickly; for inactively traded assets, business time passes slowly. Microstructure invariance hypothesizes that microstructure characteristics, which vary when measured in units of calendar time, become constants—“microstructure invariants”—when measured in units of business time.

In section 1, we formulate the three invariance principles as empirical hypotheses, conjectured to apply for all securities and across time:

- The dollar distribution of the risk transferred by a bet is the same when the risk transferred by a bet is measured in units of business time.
- The dollar transactions cost of executing a bet is the same function of the size of the bet, when size of the bet is calculated as the amount of risk transferred by the bet per unit of business time.
- Market efficiency (in volatility units) and market resilience are the same when measured in units of business time.

When measured in calendar time, the size distribution of risk transfers, the number of bets, illiquidity, bid-ask spreads, long-term market impact, efficiency, and resiliency become proportional to powers of market velocity. The velocity itself is proportional to the two-thirds power of calendar-time “trading activity,” which we define as the product of empirically observable dollar volume and volatility. This gives specific testable empirical content to the invariance hypotheses in terms of invariance relationships. For example, the size distribution of bets, as a fraction of trading volume, is inversely proportional to the two-thirds power of trading activity. The calendar-time transactions cost function is the product of an invariant cost function and asset-specific measure of illiquidity, which is proportional to the cube root of the ratio of returns variance to dollar volume. Other stock characteristics are also functions of observable dollar volume and volatility.

Given values of a tiny number of proportionality constants, the invariance relationships allow microscopic features of the market for a financial asset, like the average size of bets, to be inferred from macroscopic market characteristics such as dollar

¹This paper is based on two companion papers: the theoretical paper “Market Microstructure Invariants: Theory and Implications of Calibration” and the empirical paper “Market Microstructure Invariants: Empirical Evidence from Portfolio Transitions” (December 2011). Originally, both papers were parts of one long manuscript “Market Microstructure Invariants” (May 2011).

volume and volatility. The units in which these proportionality factors are measured are consistent with their intended economic content. Making empirical predictions on the basis of invariance principles is well established in physics. Our analysis is similar in spirit to inferring the size and number of molecules in a mole of gas from measurable large-scale physical quantities.

In section 2, we develop a meta-model showing that all three microstructure invariance hypotheses are consistent with a dynamic infinite-horizon model of market microstructure with informed trading, noise trading, intermediation (market making), and endogenous production of information. This meta-model shows that invariance relationships are ultimately related to granularity of information flow based on the underlying economics. The invariance relationships are derived under the assumption that the effort required to generate one discrete bet has distributions which do not vary across stocks and time. The invariance of market efficiency and resiliency requires an additional assumption that the signal-to-noise ratio per bet is constant across stocks and time.

In section 3, we discuss why invariance does not undermine or contradict other theoretical models of market microstructure. Instead, it builds a bridge from theoretical models to empirical tests of those models. Invariance provides guidance on what constitute good empirical proxies for some difficult-to-observe microstructure concepts such as “order imbalances.” It imposes a discipline on empirical tests by showing how to specify regressions and scale explanatory variables so that estimated regression coefficients can be assumed to be constant across observations.

In section 4, we describe the dataset of portfolio transitions data used to test invariance relationships concerning bet size and transactions costs. The dataset consists of more than 400,000 portfolio transition orders executed over the period 2001-2005 by a leading vendor of portfolio transition services. Portfolio transitions are used by institutional sponsors to transfer funds from legacy portfolio managers to new managers to replace fund managers, change asset allocations, or accommodate cash inflows and outflows. Portfolio transitions provide a good natural experiment for identifying bets and measuring transactions costs.

In section 5, we examine empirical evidence concerning the predictions of invariance for bet size, assuming that portfolio transition orders are proportional to bets. We find that the size distribution of the product of the ratio of order size as a fraction of average daily volume and the two-thirds power of trading activity indeed resembles an invariant log-normal distribution, see figure 2. Results from the regression analysis also confirm this finding.

The bets have a log-normal distribution with estimated log-variance of 2.53. The log-normal empirical distribution of bet size (a bi-modal “signed” log-normal distribution for signed bets) has much more kurtosis than the normal distribution often assumed for analytical convenience in the theoretical literature. The fat tails of the estimated log-normal distribution suggest that very large bets dominate trading volume and dominate volatility even more so. Execution of large bets may trigger noticeable

market dislocations. We suspect the log-normal empirical distribution of bets may be related to the distribution of the size of financial firms.

In section 6, we use implementation shortfall to examine whether transactions costs are consistent with the invariance hypothesis. Even though statistical tests usually reject the invariance hypotheses, their results are economically close to those implied by invariance. We find that cost functions can be closely approximated by the product of asset-specific illiquidity measure (proportional to ratio of volatility and the one-third power of trading activity) and invariant function (see figure 4). Invariance itself does not impose a particular functional form on that function, but we find empirically that it is somewhat better explained by the square root model, while a linear model better fits transactions costs for large orders in active markets (both models include spread). According to the meta-model, half of the transactions cost is, on average, due to permanent price changes and half of the transactions cost is due to temporary price deviations. We also show that quoted spreads conform reasonably closely to the predictions of invariance.

The potential benefits of invariance principles for empirical market microstructure are enormous. In the area of transactions cost measurement, for example, controlled experiments are costly and natural experiments are rare; even well-specified tests of transactions cost models tend to have low statistical power. Market microstructure invariance defines parsimonious structural relationships leading to precise predictions about how various microstructure characteristics including transactions costs vary across stocks with different dollar volume and volatility. These predictions can be tested with structural estimates of a handful of parameters using limited data from many different stocks.

In section 7, we use the estimates from our empirical tests based on the portfolio transition data to calibrate microstructure invariants and discuss implied quantitative relationships. Those implications depend on our assumptions about several additional parameters, necessary for interpreting the empirical estimates: how much volume can be attributed to trading of long-term investors rather than intermediaries, how much volatility is induced by trades rather than public announcements, and how much larger portfolio transitions are than typical bets. In the future, a better calibration and triangulation of those parameters and invariants themselves will ultimately sharpen implications of invariance hypotheses.

In both physics and market microstructure, application of invariance principles requires that certain assumptions be met. For example, the laws of physics hold in simplest form for objects traveling in a vacuum, but have to be modified when resistance from air generates friction. Similarly, in market microstructure, we believe that the invariance relationships may hold only under idealized conditions. For example, invariance relationships may assume an idealized environment with features like very small tick size, competitive market makers, and minimal transactions fees and taxes. Invariance principles provide a benchmark from which the importance of frictions such as a large tick size, non-competitive market access, or high fees and taxes can

be measured.

The idea of using invariance principles in finance and economics, at least implicitly, is not new. The theory of Modigliani and Miller (1958) is an example of an invariance principle. The idea of measuring trading in financial markets in business time or transaction time is not new either. The “time-change” literature has a long history, beginning with Mandelbrot and Taylor (1967), who link business time to transactions, and Clark (1973), who links business time to volume. More recent papers include Hasbrouck (1999), Ané and Geman (2000), Dufour and Engle (2000), Plerou et al. (2000), and Derman (2002). By applying invariance principles based on business time to market microstructure, we shift the intuition of the time-change literature from understanding the relationship between trading volume and business time to understanding the relationship between risk transfer and business time.

1 Market Microstructure Invariance as an Empirical Hypothesis

Microstructure characteristics such as order size, order arrival rate, price impact, bid-ask spread, price resilience, and market efficiency vary across assets and across time. We define “market microstructure invariance” as the empirical hypothesis that this variation almost disappears when these characteristics are examined at an asset-specific “business time” scale which measures the rate at which risk transfer takes place. Although the discussion below is based on cross-sectional implications of invariance for equity markets for individual stocks, we think that invariance principles generalize to markets for commodities, bonds, currencies, and aggregate indices such as S&P 500 futures contracts. For simplicity, we assume that a bet transfers only idiosyncratic risk about a single security, not the market risk; modeling both idiosyncratic and market risks is a subject for future research that would require developing a more complicated factor model.

In the market for an individual stock, institutional asset managers buy and sell shares to implement “bets.” We think of a bet as a decision to acquire a long-term position of a specific size in a stock, distributed approximately independently from other such decisions. Intermediaries with short-term trading strategies—market makers, high frequency traders, and other arbitragers—clear markets by taking the other side of bets placed by long-term traders.

Notation. Over short periods of time, we assume that the bet arrival rate can be approximated by a compound Poisson process, with γ denoting the bet arrival rate of independently distributed bets, measured in units of bets per calendar day, and \tilde{Q} denoting a random variable with probability distribution representing the signed size of bets, measured in shares (positive for buys, negative for sells), where $E\{\tilde{Q}\}$ is

approximately zero. The bet arrival rate γ measures market velocity, the rate at which business time passes for a particular stock.

Over long periods of time, we assume that the inventories of intermediaries do not grow in an unbounded manner; this requires bets to have small negative autocorrelation. Furthermore, both the bet arrival rate and the distribution of bet size change over longer periods of time as the level of trading activity in a stock increases or decreases.

Bets can be difficult for researchers to observe. Consider an asset manager who places one bet by purchasing 100,000 shares of IBM stock. The bet might be implemented by placing orders over several days, and each of the orders might be “shredded” into many small trades showing up on the ticker at various prices. Since bets represent independent increments in the intended order flow, the various trades which implement the bet should all be added together to recover the size of the original bet. Bets may be difficult to identify from TAQ data.

Similarly, if an analyst issues a buy recommendation to ten different customers and each of the customers quickly places executable orders to buy 10,000 shares, it might be appropriate to think of the ten orders as one bet for 100,000 shares. The ten individual orders lack statistical independence. The bet results from a new idea, which can be shared.

We assume that, on average, each unit of bet volume results in ζ units of total volume, i.e., one unit of bet volume leads to $\zeta - 1$ units of intermediation volume. On a given calendar day, expected trading volume (in shares) is given by $V := \zeta / 2 \cdot \gamma \cdot E|\tilde{Q}|$ (counting a buy matched to a sell only once). We define “expected bet volume” by

$$\bar{V} := \gamma \cdot E|\tilde{Q}| = \frac{2}{\zeta} \cdot V. \quad (1)$$

We can estimate expected bet volume \bar{V} by combining an estimate of expected market volume V with a value for the “intermediation multiplier” ζ . If all trades are bets and there are no intermediaries, then $\zeta = 1$, since each unit of trading volume would match a buy-bet with a sell-bet. If a monopolistic specialist intermediates all bets without involvement of other intermediaries, then $\zeta = 2$. If each bet is intermediated by different market makers, each of whom lays off inventory by trading with other market makers, then $\zeta = 3$. If positions are passed around among multiple intermediaries, then $\zeta \geq 4$.

Let σ denote the percentage standard deviation of a stock’s daily returns. Some price fluctuations result from release of information directly without trading, such as overnight news announcements. Let ψ^2 denote the fraction of returns variance σ^2 which results from order flow imbalances, which we assume ultimately result from bets. We define “trading volatility” as the standard deviation of returns resulting from bet-related order flow imbalances:

$$\bar{\sigma} := \psi \cdot \sigma. \quad (2)$$

Let P denote the price of the stock; then dollar trading volatility is $P \cdot \bar{\sigma} = \psi \cdot P \cdot \sigma$.

Invariance of Bets. In one unit of business time $1/\gamma$, a bet of dollar size $P\tilde{Q}$ generates a standard deviation of dollar mark-to-market gains or losses equal to $P|\tilde{Q}| \cdot \bar{\sigma}\gamma^{-1/2}$. The signed standard deviation, $P\tilde{Q} \cdot \bar{\sigma}\gamma^{-1/2}$, which is positive for buys and negative for sells, measures both the direction and the size of the risk transfer resulting from the bet. It is measured in dollars per unit of business time $1/\gamma$.

Market microstructure invariance hypothesizes that *the dollar distribution of risks transferred by bets is the same for all stocks when the risk transferred by a bet is measured in units of business time*. Since $P\tilde{Q} \cdot \bar{\sigma}\gamma^{-1/2}$ measures the risk transferred by a bet per unit of business time, invariance implies that the distribution of $P\tilde{Q} \cdot \bar{\sigma}\gamma^{-1/2}$ does not vary across stocks. Letting “ \sim ” mean “is equal in distribution to,” there is some random variable \tilde{I} with an “invariant distribution” such that for all stocks,

$$P\tilde{Q} \cdot \bar{\sigma}\gamma^{-1/2} \sim \tilde{I}. \quad (3)$$

The distribution of risk transfer \tilde{I} is a market microstructure invariant.

By analogy with bets, we define “trading activity” W as the product of expected dollar trading volume PV and calendar returns volatility σ , i.e., $W := \sigma \cdot PV$. Similarly, define “bet activity” \bar{W} as the product of dollar bet volume $P\bar{V}$ and trading volatility $\bar{\sigma}$, i.e., $\bar{W} := \bar{\sigma} \cdot P\bar{V}$. Given values of the volume multiplier ζ and the trading volatility factor ψ , we can convert more-easily-observed trading activity W into less-easily-observed bet activity \bar{W} using the relationship $\bar{W} = W \cdot 2\psi/\zeta$.

Since equation (3) implies $\tilde{Q} \sim \gamma^{1/2}P^{-1}\bar{\sigma}^{-1} \cdot \tilde{I}$ and expected bet volume $\bar{V} = \gamma \cdot E|\tilde{Q}|$, bet activity \bar{W} can be expressed as a function of the unobservable speed of business time γ :

$$\bar{W} = \bar{\sigma} \cdot P \cdot \bar{V} = \bar{\sigma} \cdot P \cdot \gamma \cdot E|\tilde{Q}| = \gamma^{3/2} \cdot E|\tilde{I}|. \quad (4)$$

In equation (4), the exponent $3/2$ has simple intuition. Suppose business time γ speeds up by a factor of 4, but calendar trading volatility $\bar{\sigma}$ does not change. Then trading volatility in units of business time $\bar{\sigma}\gamma^{-1/2}$ falls by $1/2$. The invariance principle (3) therefore requires bet size \tilde{Q} to increase by a factor of 2 to keep the distribution of \tilde{I} invariant. The resulting increase in bet volume by a factor of $8 = 4^{3/2}$ can be decomposed into an increase in the number of bets by a factor of $8^{2/3} = 4$ and the size of bets by a factor of $8^{1/3} = 2$. As trading activity increases, the number of bets increases twice as fast as their size.

Invariance makes it possible to infer the bet arrival rate γ and the average size of bets $E|\tilde{Q}|$ from the level of bet activity \bar{W} , up to some proportionality constant which does not vary across stocks. Define the constant $\iota := (E|\tilde{I}|)^{-1/3}$. Solving equation (4) for γ in terms of \bar{W} yields

$$\gamma = \bar{W}^{2/3} \cdot \iota^2, \quad E|\tilde{Q}| = \bar{W}^{1/3} \cdot \frac{1}{P\bar{\sigma}} \cdot \iota^{-2}. \quad (5)$$

The shape of the entire distribution of bet size \tilde{Q} can be obtained by plugging γ from

(5) into (3). Expressing bet size \tilde{Q} as a fraction of expected bet volume \bar{V} , we obtain

$$\frac{\tilde{Q}}{\bar{V}} \sim \bar{W}^{-2/3} \cdot \tilde{I} \cdot \iota. \quad (6)$$

Equations (5)) and (6) summarize the implications of invariance for bet size and arrival rate. We test these implications in section 5 below.

Invariance of Transactions Costs. Market microstructure invariance also makes empirical predictions about transactions costs. Market microstructure invariance hypothesizes that *the dollar expected transactions cost of executing a bet is the same function of the size of the bet for all stocks, when the size of the bet is calculated as the dollar amount of risk transferred by the bet per unit of business time.* Since the risk transferred per unit of business time by a bet of \tilde{Q} shares is measured by $\tilde{I} = P\tilde{Q} \cdot \bar{\sigma}\gamma^{-1/2}$, invariance of trading costs implies that there is an “invariant transactions cost function” $C_B(\tilde{I})$ which measures the execution cost of transferring the risk represented by $\tilde{Q} = \tilde{I}/(\bar{\sigma}P\gamma^{-1/2})$ shares. The transactions cost function $C_B(\cdot)$ is a market microstructure invariant.

Suppose, for example, that a 99th percentile bet in stock *A* is for \$10million (e.g., 100,000 shares at \$100 per share) while a 99th percentile bet in stock *B* is for \$1 million (e.g., 100,000 shares at \$10 per share). The invariance of the distribution of bet size implies that value of \tilde{I} is the same for both bets because they occupy the same percentile in the bet size distribution for their respective stocks. Even though the bet in stock *A* has 10 times the dollar value of the bet in stock *B*, invariance of transactions costs implies that the expected dollar cost of executing each bet is the same because both bets transfer the same amount of risk per stock-specific unit of business time. The function $C_B(\cdot)$ is the same function for all stocks. Measured in basis points, however, invariance implies that the transactions cost for Stock *B* is 10 times greater than for stock *A*.

Let $C(\tilde{Q})$ denote the stock-specific cost of executing a bet of \tilde{Q} shares, expressed as a fraction of the notional value of the bet $|P\tilde{Q}|$ (i.e., in units of 10^{-4} basis points). Define $\bar{C}_B := E\{C_B(\tilde{I})\}$. Using equation (3) yields

$$C(\tilde{Q}) = \frac{C_B(\tilde{I})}{|P\tilde{Q}|} = \frac{\bar{C}_B}{E|P\tilde{Q}|} \cdot \frac{C_B(\tilde{I})/\bar{C}_B}{|\tilde{I}|/E|\tilde{I}|}. \quad (7)$$

Let $f(\tilde{I}) := [C_B(\tilde{I})/\bar{C}_B]/[|\tilde{I}|/E|\tilde{I}|]$ denote the invariant “average cost function” for executing a bet \tilde{I} . This function is defined in terms of deviations of functions $C_B(\tilde{I})$ and $|\tilde{I}|$ from their means. For example, if I denotes a bet 5 times greater than than mean bet size $E|\tilde{I}|$ and such a bet has a transactions cost 10 times greater than the mean cost \bar{C}_B , then $f(I) = 2$.

Let $1/L := \bar{C}_B/E|P\tilde{Q}|$ be an asset-specific measure of illiquidity equal to the dollar-volume-weighted expected cost of executing a bet. For an asset manager who

places many bets in the same stock, this expresses expected transactions cost as a fraction of the dollar value traded (basis points $\times 10^{-4}$). Equation (5) yields (recall $\iota := (E|\tilde{I}|)^{-1/3}$)

$$1/L := \bar{\sigma}\bar{W}^{-1/3} \cdot \iota^2\bar{C}_B = [P\bar{V}/\bar{\sigma}^2]^{-1/3} \cdot \iota^2\bar{C}_B. \quad (8)$$

The cost of executing a bet of \tilde{Q} shares can be written

$$C(\tilde{Q}) = \bar{\sigma}\bar{W}^{-1/3} \cdot \iota^2\bar{C}_B \cdot f\left(\frac{\bar{W}^{2/3}}{\iota} \cdot \frac{\tilde{Q}}{\bar{V}}\right) = \frac{1}{L} \cdot f(\tilde{I}). \quad (9)$$

The cost function is the product of the asset-specific illiquidity measure $1/L$ and an invariant transactions cost function $f(\tilde{I})$. We test this relationship empirically in section 6.

The liquidity measure $L = [\iota^2\bar{C}_B]^{-1} \cdot [P\bar{V}/\bar{\sigma}^2]^{1/3}$ is an intuitive and practical alternative to other measures of liquidity, such as Amihud (2002) and Stambough and Pastor (2003). To implement L empirically, it is simpler to define L in terms of expected dollar trading volume PV and expected returns volatility σ rather than in terms of dollar bet volume $P\bar{V}$ and trading volatility $\bar{\sigma}$. We have

$$L = \left[\frac{2(\iota^2\bar{C}_B)^{-3}}{\zeta\psi^2}\right]^{1/3} \cdot \left[\frac{PV}{\sigma^2}\right]^{1/3}. \quad (10)$$

The idea that liquidity is related to dollar volume per unit of returns variance PV/σ^2 is intuitive. Traders believe that transactions costs are low in markets with high dollar volume and high in markets with high volatility. If the intermediation multiplier ζ and trading volatility factor ψ do not vary across stocks, then $L \propto [PV/\sigma^2]^{1/3}$ becomes a simple index of liquidity.

The liquidity measure $L = (\iota^2\bar{C}_B)^{-1} \cdot [P\bar{V}]^{1/3} \cdot [\bar{\sigma}]^{-2/3}$ is similar to the definition of “market temperature” $\chi = \bar{\sigma} \cdot \gamma^{1/2}$ in Derman (2002); substituting for γ from equation (5), we obtain $\chi = \iota \cdot [P\bar{V}]^{1/3} \cdot [\bar{\sigma}]^{4/3} \propto L \cdot \sigma^2$.

Invariance does not imply a specific functional form for $f(\cdot)$. In our analysis, we focus on two specific functional forms: linear price impact costs and square root price impact costs. For both functional forms, we also allow a constant bid-ask spread cost component. Linear price impact is consistent with price impact models based on adverse selection, such as Kyle (1985). Square root price impact functions are consistent with empirical findings in the econophysics literature, such as Gabaix et al. (2006); some papers in this literature find an exponent closer to 0.60 than the square root exponent 0.5, such as Almgren et al. (2005).

For the linear model, we write $f(\tilde{I})$ as the sum of a bid-ask spread component and a linear price impact cost component, $f(\tilde{I}) := (\iota^2\bar{C}_B)^{-1} \cdot \bar{\kappa}_0 + (\iota\bar{C}_B)^{-1} \cdot \bar{\kappa}_I \cdot |\tilde{I}|$, where invariance implies that the bid-ask spread cost parameter $\bar{\kappa}_0$ and the market impact cost parameter $\bar{\kappa}_I$, as well as constants ι and \bar{C}_B , do not vary across stocks.

The linearity of $f(\cdot)$ as a function of $|\tilde{I}|$ implies that $C_B(\tilde{I})$ is a quadratic function of $|\tilde{I}|$. The proportional cost function $C(\tilde{Q})$ from (9) is therefore given by

$$C(\tilde{Q}) = \bar{\sigma} \left[\bar{\kappa}_0 \cdot \bar{W}^{-1/3} + \bar{\kappa}_I \cdot \bar{W}^{1/3} \cdot \frac{|\tilde{Q}|}{\bar{V}} \right]. \quad (11)$$

When bets are measured as a fraction of expected trading volume and transactions costs are measured in basis points, bid ask spread costs are decreasing in bet activity \bar{W} and market impact costs are increasing in bets activity \bar{W} . When transactions costs in basis points are further scaled in units of trading volatility $\bar{\sigma}$, equation (11) says that bid-ask spread costs are proportional $\bar{W}^{-1/3}$ and market impact costs are proportional to $\bar{W}^{1/3}$ for a given fraction of volume.

For the square root model, we write $f(\tilde{I})$ as the sum of a bid-ask spread component and a square root function of $|\tilde{I}|$, obtaining $f(\tilde{I}) := (\iota^2 \bar{C}_B)^{-1} \bar{\kappa}_0 + (\iota^{3/2} \bar{C}_B)^{-1} \bar{\kappa}_I \cdot |\tilde{I}|^{1/2}$, where invariance implies that $\bar{\kappa}_0$, $\bar{\kappa}_I$, ι , and \bar{C}_B do not vary across stocks. The proportional cost function $C(\tilde{Q})$ from (9) is given by

$$C(\tilde{Q}) = \bar{\sigma} \left[\bar{\kappa}_0 \cdot \bar{W}^{-1/3} + \bar{\kappa}_I \cdot \left| \frac{\tilde{Q}}{\bar{V}} \right|^{1/2} \right]. \quad (12)$$

When transactions costs are measured in units of trading volatility $\bar{\sigma}$, bid-ask spread costs remain proportional to $\bar{W}^{-1/3}$, but the square root model implies that the bet activity coefficient $\bar{W}^{1/3}$ cancels out of the price impact term. Indeed, the square root is the only function for which invariance leads to the empirical prediction that impact costs (measured in units of returns volatility) depend only on bet size as a fraction of bet activity \tilde{Q}/\bar{V} and are not a function of any other stock characteristics, including the level of bet activity \bar{W} . In the context of invariance, the square root model places the strongest possible empirically testable restrictions on which characteristics of the market for a stock can affect transactions costs. If there are no bid-ask spread costs ($\bar{\kappa}_0 = 0$), then the square root model implies the parsimonious transactions cost function $C(\tilde{Q}) = \bar{\sigma} \cdot \bar{\kappa}_I \cdot [|\tilde{Q}|/\bar{V}]^{1/2}$. We calibrate cost functions in section 5 below.

The model in section 2 will show that permanent and temporary components of transactions costs each accounts for exactly half of total costs. These components do not necessary correspond directly to a linear (or square root) terms and fixed bid-ask spread terms in $C(\tilde{Q})$. Some price impact costs may be temporary.

Invariance of Market Efficiency and Resilience. Black (1986) defines an efficient market as “one in which price is within a factor 2 of value.” We think of “fundamental value” F as the value to which a stock price would converge if traders continuously expended huge resources acquiring information about its value. Let Σ denote the variance of the log-difference between price and fundamental value: $\Sigma := \text{var}\{\ln(P/F)\}$. We measure “market efficiency” by $1/\Sigma^{1/2}$. If the “factor of 2”

in Black's definition represents a ξ standard deviation event, then the spirit of Black's definition implies that a market is efficient if $1/\Sigma^{1/2} \leq \xi^{-1} \cdot \ln(2)$.

Our measure of market resilience is the mean-reversion parameter ρ (per calendar day) measuring the speed with which a random shock to prices, resulting from execution of an uninformative bet, dies out over time. The half-life of an uninformative shock to prices is $\rho^{-1} \cdot \ln(2)$.

Market microstructure invariance hypothesizes that (1) *market efficiency is the same for all stocks if measured in units of volatility per unit of business time* and (2) *market resilience is the same for all stocks if measured in units of business time*. Invariance of market efficiency implies that the ratio $\Sigma^{1/2}/[\bar{\sigma}\gamma^{-1/2}]$ is invariant across stocks. Since $L \propto \bar{\sigma}\gamma^{-1/2}$, invariance implies proportionality between liquidity L and efficiency $1/\Sigma^{1/2}$. Invariance implies that the ratio ρ/γ is invariant across stocks. Thus, invariance implies that resilience is proportional to market velocity, the rate at which business time passes. Using equation (5), invariance implies (recall $\iota := (E|\tilde{I}|)^{-1/3}$)

$$\Sigma^{1/2} \sim \bar{\sigma} \cdot \bar{W}^{-1/3} \cdot \iota^{-1}. \quad (13)$$

$$\rho \sim \bar{W}^{2/3} \cdot \iota^2. \quad (14)$$

When trading activity increases by a factor of 8, invariance implies that resilience increases by a factor of 4 and market efficiency increases by a factor of 2. Invariance suggests that the factor of 2 in the definition of market efficiency in Black (1986) should be modified to vary across stocks according to equation (13).

Intuitively, the unstated invariant proportionality factors implied by equations (13) and (14) should be related to the information content of bets. More informative bets should make markets more efficient and resilient. This intuition is made precise in the meta-model in section 2.

We do not examine empirically the predictions of equations (13) and (14); they are interesting topics for future research.

Discussion. Invariance implies that trading liquidity and funding liquidity may be two sides of the same coin. Trading liquidity is measured by $L \propto \bar{\sigma} \cdot \bar{W} \propto \bar{\sigma} \cdot \gamma^{-1/2}$. A good measure of funding liquidity is the repo haircut that sufficiently protects a creditor from losses if the creditor sells the collateral due to default by the borrower. Such a haircut should be proportional to the volatility of the asset's return over the horizon during which the collateral would be liquidated. Invariance of resiliency ρ suggests that this horizon should be proportional to business time $1/\gamma$, making volatility over the liquidation horizon proportional to $\bar{\sigma} \cdot \gamma^{-1/2}$, which is proportional to L . Thus, both trading liquidity and funding liquidity are measured by L .

The velocity of the market suggests a speed with which collateral should be liquidated without disrupting the normal price-formation process. In a fire sale, collateral is liquidated very quickly relative to the natural velocity of the market, leading to short-term over-reaction and high liquidation costs.

Invariance is consistent with the Modigliani-Miller irrelevance of leverage and splits. Invariance relationships do not change if a company levers up its equity by paying a debt-financed cash dividend or implements a stock split.

Invariance is also consistent with irrelevance of the units in which time is measured. This is unlike some other models, such as ARCH and GARCH. The values of \tilde{I} , $C_B(\tilde{I})$, $f(\tilde{I})$, and $1/L$ —and therefore the economic content of the predictions of invariance—remain the same regardless of whether researchers measure γ , \bar{V} , $\bar{\sigma}$, and $\bar{W}^{2/3}$ using daily weekly, monthly, or annual units of time.

The values of \tilde{I} and $C_B(\tilde{I})$ are measured in dollars. If invariance relationships are applied to an international context in which markets have different currencies or different real exchange rates or applied across periods of time where the price level is changing significantly, invariance is consistent with the idea that these nominal values should be deflated by the real productivity-adjusted wages of finance professionals in the local currency of the given market. Like fundamental constants in physics, such deflation would make the invariants \tilde{I} and $C_B(\tilde{I})$ dimensionless.

We do not expect invariance to hold perfectly across different markets and different times periods. We expect transactions costs, particularly bid-ask spread costs, to be influenced by numerous institutional features, such as government regulation (e.g., short sale restrictions or customer order handling rules), transactions taxes, competitiveness of market making institutions and trading platforms, tick size, market fragmentation, and technological change. To the extent that, say, minimum tick size rules affect bid-ask spread costs, we believe that market microstructure invariance can be used as a benchmark against which the effect of tick size on bid-ask spread costs can be evaluated.

2 Market Microstructure Invariance as an Implication of a Structural Meta-Model

In this section, we derive invariance relationships as endogenous implications of a steady-state structural “meta-model” of informed trading, noise trading, and intermediation (market making).

Our set-up has the following structure. The unobserved “fundamental value” of the stock follows geometric Brownian motion with log-standard-deviation σ . Informed traders face given costs c_I of acquiring information of given precision τ ; they place informed bets \tilde{Q} which incorporate a given fraction θ of the information into prices. Noise traders place bets which turn over a constant fraction η of the stock’s float of N shares, mimicking the trading of informed traders even though their private “signal” has no information value, as in Black (1986). Intermediaries set prices by filtering the order flow for information about the fundamental value. They lose money from being “run over” by informed bets, but they break even from bid-ask spread costs, temporary impact costs, or other trading costs imposed on all traders. The model

endogenously determines the rate of informed trading γ_I , the rate of uninformed trading γ_U , the distribution of bet sizes \tilde{Q} , market efficiency $1/\Sigma^{1/2}$, market resilience ρ , the illiquidity measure $1/L$, and a long-term permanent price impact parameter λ which in the long run reveals the information content of the order flow.

Invariance relationships come about through the following intuition: Suppose the number of noise traders increases for some exogenous reason. In the meta-model, this happens when market capitalization increases, keeping the share turnover of noise traders constant. As a result, market depth increases and, consequently, the number of informed traders increases, since their bets now are more profitable. If the number of informed traders increases by a factor of 4, then each of their bets accounts for a 4 times smaller fraction of returns variance. The volatility per unit of business time decreases by a factor of 2. The meta-model shows that market efficiency and liquidity both increase by a factor of 2, as a result of which informed traders exactly cover the cost of private signals by submitting bets 2 times as large as before. The overall dollar volume in the market increases by a factor of 8. As a result, the “one-third, two-thirds” intuition comes about: One-third of the increase in dollar volume comes from changes in bet size ($8^{1/3} = 2$) and two-thirds comes from changes in the number of bets ($8^{2/3} = 4$).

We call our framework a “meta-model” because, unlike Kyle (1985), we do not model explicitly the process by which informed traders dynamically execute bets and intermediaries dynamically set prices in continuous time. Our meta-model becomes a closed “model” with the same invariance properties when we make the explicit simplifying assumption that informed and noise traders sequentially enter the market and trade only once, at one price, as in Glosten and Milgrom (1985).

Although the model is motivated by the time series properties of a single stock as its market capitalization changes, the model applies cross-sectionally across different stocks under the assumption that the exogenously assumed cost of a private signal c_I is constant across all stocks. We show that this one “deep” structural assumption imposes a granularity on information which drives invariance relationships based on the granularity of bet size.

Both the invariance of bets and the invariance of trading costs hold precisely when the cost of a signal c_I is constant in the form hypothesized in section 1 when volatility σ , float N , noise turnover rate η , the fraction of informed bets (which we show equals θ), the precision of informed signals τ , and share price P vary across stocks. The model reveals that the invariance of market efficiency and resilience requires stronger assumptions: The informativeness of a bet, measured as the product of signal precision τ and the squared fraction of informed traders θ^2 , must be constant across stocks.

In the remainder of this section, we sketch out the details of the meta-model, using notation consistent with the previous section. The meta-model operates with concepts of bet volume \bar{V} and bet volatility $\bar{\sigma}$; for notational convenience we assume $V = \bar{V}$ and $\sigma = \bar{\sigma}$. It is straightforward to adjust the meta-model by applying

equations (1) and (2).

Fundamental Value and Private Information. Let the unobserved “fundamental value” of the asset follow a geometric brownian motion given by $V(t) := \exp[\sigma \cdot B(t) - \sigma^2 t/2]$, where $B(t)$ follows standardized Brownian motion (with $\text{var}\{B(t + \Delta t) - B(t)\} = \Delta t$) and the constant σ measures the volatility of fundamental value. Based on the history of past order flow, we assume that the market’s conditional estimate of $\sigma \cdot B(t)$ is distributed approximately $N[\sigma \cdot \bar{B}(t), \Sigma(t)]$. This is consistent with the price being given approximately by $P(t) = \exp[\sigma \bar{B}(t) + \Sigma(t)/2 - \sigma^2 t/2]$. Here $\Sigma^{1/2}$ measures the standard deviation of the log-difference between price and fundamental value; $1/\Sigma^{1/2}$ measures market efficiency.

When the n th bet is informed, the informed trader observes at some date t_n a signal \tilde{i}_n given by $\tilde{i}_n = \tau^{1/2} \cdot \Sigma^{-1/2}(t_n) \cdot \sigma \cdot [B(t_n) - \bar{B}(t_n)] + \tilde{Z}_{In}$, where τ is an exogenous constant parameter measuring the precision of the signal and the noise $\tilde{Z}_{In} \sim NID(0, 1)$ is distributed independently from the process $B(t)$. Note that τ also measures the signal-to-noise ratio. We assume τ is small enough that $\text{var}\{\tilde{i}_{In}\} \approx 1$. For notational convenience, we suppress the subscripts n from here on. When an informed trader observes a signal \tilde{i} , he updates his estimate of $B(t)$ from $\bar{B}(t)$ to $\bar{B}(t) + \Delta B_I(t)$. Using a continuous time linear approximation in which τ is small, we have

$$\Delta B_I(t) \approx \tau^{1/2} \cdot \Sigma^{1/2} / \sigma \cdot \tilde{i}(t). \quad (15)$$

The dollar price change implied by $\Delta B_I(t)$ is approximately

$$E\{V(t) - P(t) \mid P(t), \Delta B_I(t)\} \approx P(t) \cdot (\exp[\sigma \cdot (\Delta B_I(t) + \Delta B_I(t)^2/2)] - 1) \approx P(t) \cdot \sigma \cdot \Delta B_I(t). \quad (16)$$

To simplify matters, we assume that filtering is linear, implying a linear long-term price impact. Even though a proper filtering—involving a mixture of normals—is not exactly linear, we conjecture a linear approximation can be used in a large-market limit in which there are many bets, each of which has small information content. We use continuous-time approximations related to steady state behavior to keep the model simple and intuitive, even though linearity holds precisely only in the limit, when dollar trading volume becomes infinite.

Informed Trading. Informed traders arrive in the market at endogenously determined rate γ_I , each informed trader acquires one private signal \tilde{i} , then places one and only one bet, which is executed by trading in some un-modeled manner over time. If we modeled the informed trader’s trading strategy explicitly, we would be describing a “model,” not a “meta-model.” Without solving an optimization problem explicitly, we assume that an informed trader executes a bet of \tilde{Q} shares as linear multiple of $\Delta B_I(t)$ in such a way that the expected long-term permanent price impact is an exogenous constant fraction θ of the impact $P(t) \cdot \sigma \cdot \Delta B_I(t)$ that would fully incorporate

the signal value into prices, i.e.,

$$\tilde{Q} = \theta/\lambda \cdot P(t) \cdot \sigma \cdot \Delta B_I(t). \quad (17)$$

If the informed trader were to incur no trading costs, his expected “paper trading” trading profits, denoted $\bar{\pi}_I$, would be

$$\bar{\pi}_I = \frac{\theta \cdot P^2 \cdot \sigma^2 \cdot E\{\Delta B_I^2\}}{\lambda}. \quad (18)$$

Expected “permanent” price impact costs from moving continuously along a linear demand schedule of slope λ , denoted C_P , are

$$\bar{C}_P = \frac{1}{2}\lambda \cdot E\{\tilde{Q}^2\} = \frac{\theta^2 \cdot P^2 \cdot \sigma^2 \cdot E\{\Delta B_I^2\}}{2 \cdot \lambda}. \quad (19)$$

Constant Rate of Noise Trading. Noise traders arrive randomly in the market at endogenous rate γ_U . Each noise trader places one bet which mimics the size distribution and unmodeled execution strategy of an informed bet even though it contains no information. Noise traders are assumed to trade randomly, turning over on average a constant percentage η of the market capitalization of the firm per day. Let informed trades be distributed as the random variable \tilde{Q} as in equation (17). If the price of the stock is P and shares outstanding is N , then market cap is $P \cdot N$ dollars, share volume from noise traders is expected to be $\eta \cdot P \cdot N$ dollars per day, and the arrival rate of noise trades γ_U solves the equation

$$\gamma_U \cdot E\{|\tilde{Q}|\} = \eta \cdot N. \quad (20)$$

The combined rate at which bets are placed by informed traders and noise traders is $\gamma = \gamma_I + \gamma_U$. Bets of informed and noise traders add up into daily trading volume,

$$\gamma \cdot E\{|\tilde{Q}|\} = V. \quad (21)$$

Permanent Market Depth. Risk neutral intermediaries (market makers) are assumed to set prices such that the permanent price impact of anonymous trades by informed and noise traders reveals on average the information in the order flow. Markets makers update prices by $\lambda \cdot \tilde{Q}$, taking into account that a bet can be either an informed bet $\tilde{Q} = \theta/\lambda \cdot P(t) \cdot \sigma \cdot \Delta B_I(t)$ with information content $P(t) \cdot \sigma \cdot \Delta B_I(t)$ and probability $\gamma_I/(\gamma_I + \gamma_U)$ or a noise bet with the same probability distribution but with no information content. The resulting linear regression coefficient is $\lambda = \gamma_I/(\gamma_I + \gamma_U) \cdot \lambda/\theta$. Canceling λ from both sides, the regression coefficient implies that arrival rates of informed and noise bets γ_I and γ_U adjust endogenously so that the probability that the bet is informed equals the exogenously assumed fraction of the informed trader’s signal incorporated into prices θ :

$$\theta = \gamma_I/(\gamma_I + \gamma_U). \quad (22)$$

Equations (20), (21), and (22) imply that, in terms of exogenous variables, V is given by $V = \eta \cdot N/(1 - \theta)$.

Temporary and Permanent Price Impact Costs. The long-term impact of a bet of size \tilde{Q} moves the price from P to $P + \lambda \cdot \tilde{Q}$. If the bet is executed by moving continuously along a linear demand schedule with slope λ , then the average execution price is $P + \lambda\tilde{Q}/2$. On average, the permanent price impact cost is $\bar{C}_P := \lambda E\{\tilde{Q}^2\}/2$. Each bet also incurs an additional “transitory” execution cost with expected value \bar{C}_T . These costs, which might represent bid-ask spread costs or temporary price impact costs, are profits for market makers. The total expected costs of executing a bet are denoted $\bar{C}_B := \bar{C}_P + \bar{C}_T$.

The equilibrium level of costs allows market makers to break even. Thus, \bar{C}_T is determined by equating the expected permanent market impact costs \bar{C}_P and other costs \bar{C}_T of both informed and uninformed traders to the expected pre-impact profits of informed traders:

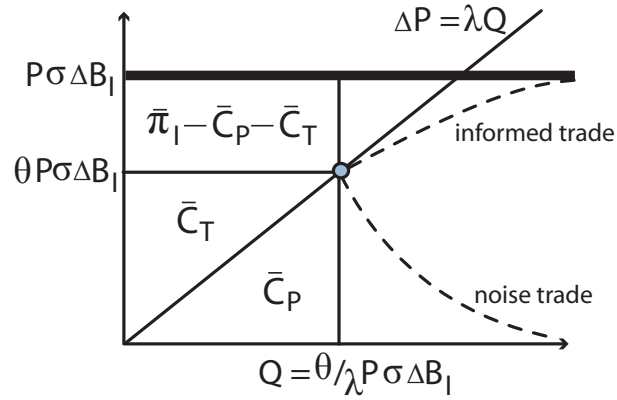
$$(\gamma_I + \gamma_U) \cdot (\bar{C}_P + \bar{C}_T) = \gamma_I \cdot \bar{\pi}_I. \quad (23)$$

Using (18) and (19), this implies

$$\bar{C}_P = \bar{C}_T = \frac{\bar{C}_B}{2} = \frac{\lambda \cdot E\{\tilde{Q}^2\}}{2} = \frac{\theta^2 \cdot P^2 \cdot \sigma^2 \cdot E\{\Delta B_I(t)^2\}}{2 \cdot \lambda}. \quad (24)$$

Alternatively, if traders announce entire quantities they want to trade and market makers set one price at which all quantities are traded, the execution price is the permanent impact price $P + \lambda\tilde{Q}$, not $P + \lambda\tilde{Q}/2$; this implies $\bar{C}_P = \bar{C}_T = \bar{C}_B/2$ and leads to the same results.

Figure 1: Intuition of Meta-Model.



There is price continuation after an informed trade and mean reversion after a noise trade.

Figure 1 illustrates informally and non-rigorously the intuition of what happens “on average.” Informed traders incorporate only fraction θ of their information into

prices, pay transactions costs $\bar{C}_P + \bar{C}_T$ and expect to make $\bar{\pi}_I - \bar{C}_P - \bar{C}_T$ in net trading profits after prices fully incorporate their information. Noise traders execute orders which would earn nothing if there were no transactions costs but incur transactions costs $\bar{C}_P + \bar{C}_T$. As in Treynor (1995), losses of market makers on trading with informed traders $\gamma_I \cdot (\bar{\pi}_I - \bar{C}_P - \bar{C}_T)$ are equal to their gains on trading with noise traders $\gamma_U \cdot (\bar{C}_P + \bar{C}_T)$.

The rate at which informed traders place bets γ_I is obtained by equating the expected profits from trading on a signal to the sum of (1) permanent market impact costs C_P , (2) other trading costs C_T , and (3) the cost of acquiring private information denoted c_I :

$$\bar{\pi}_I = \bar{C}_P + \bar{C}_T + c_I. \quad (25)$$

It is necessary to have temporary impact costs in the meta-model to sustain an equilibrium. Without temporary impact costs, it would be impossible for market makers to break even trading at the average price of $P + \lambda\tilde{Q}/2$ instead of the break-even “permanent impact” price of $P + \lambda\tilde{Q}$. Because of temporary impact, the “last” trades in a bet will usually be executed at prices higher than the long-term permanent impact value of $P + \lambda\tilde{Q}$. Market makers will usually be making profits on the “last” trades in a bet during the subsequent mean-reversion of prices to the long-term level, even though they will usually be losing money from the very “first” trades in a bet during the subsequent continuation of prices to the long-term level. On average, market makers break even. Somewhat similar intuition underlies a “fair-pricing” rule saying that the average execution price has to be equal to the post-trade reversion price, as suggested in Farmer et al. (2012). Our use of the terms “permanent” and “transitory” are somewhat non-standard.

The existence of temporary impact is important not only in the meta-model but also in most models, in which large bets are executed as sequences of many small trades. In the continuous-time version of Kyle (1985), an informed trader executes many “small” positively correlated trades X of order dt at price increments $\lambda \cdot X$ of order dt ; the informed trader’s “transitory” bid-ask spread costs of order dt^2 are economically inconsequential. Noise traders dominate volume with “large” trades X of order $d\tilde{B}$ at price increments $\lambda \cdot X$ of order $d\tilde{B}$, continuously paying to market makers transitory bid-ask spread costs of order dt . This allows market makers to break even, even though they make losses trading with informed traders.

As in section 1, let $1/L := \bar{C}_B/E|P\tilde{Q}|$ denote the expected cost of executing a bet, denoted in basis points.

Kalman Filter. In a steady state (which is reached only as an approximation), the volatility of prices reflects the arrival of new information, implying

$$\gamma \cdot \theta^2 \cdot P^2 \cdot \sigma^2 \cdot \text{var}\{\Delta B_I(t)\} = P^2 \cdot \sigma^2, \quad (26)$$

Note that $\theta^2 \cdot \sigma^2 \cdot \text{var}\{\Delta B_I(t)\}$ measures returns variance per unit of business time while σ^2 measures returns variance per unit of calendar time. After canceling $P \cdot \sigma$,

it follows that

$$\text{var}\{\Delta B_I(t)\} = \frac{1}{\gamma \cdot \theta^2}. \quad (27)$$

Since $\Delta B_I(t) \approx \tau^{1/2} \cdot \Sigma^{1/2} / \sigma \cdot \tilde{i}(t)$ and $\text{var}(\tilde{i}) \approx 1$, we have

$$\Sigma / \sigma^2 = \frac{1}{\gamma \cdot \tau \theta^2}. \quad (28)$$

Since $\Sigma = \sigma^2 \cdot \text{var}\{B(t) - \bar{B}(t)\}$, the value of $1/\Sigma^{1/2}$ measures “market efficiency” consistently with section 1, as the accuracy with which market prices reveal the unobserved fundamental value of the asset. More accurate prices reduce the value of private signals, in this sense making it harder for informed traders to profit from their private information. According to equation (28), more accurate signals (increasing τ) and more frequent bets (increasing γ) make market price more efficient in a steady state (reducing $\Sigma^{1/2}$).

Market efficiency is closely related to resiliency. As a result of each bet, market makers update their estimate of $B(t) - \bar{B}(t)$. A trade is informed with probability θ and, if informed, incorporates a fraction θ of its information content into prices, leading to a price update $\theta \tau^{1/2} \Sigma^{1/2} / \sigma \cdot \{\tau^{1/2} \Sigma^{-1/2} \cdot \sigma \cdot [B(t) - \bar{B}(t)] + \tilde{Z}_I\}$ from (15). A trade is uninformed with probability $1 - \theta$, adding noise $\theta \cdot \tau^{1/2} \Sigma^{1/2} \cdot \sigma \cdot \tilde{Z}_U$ into prices. As a result, the error $B(t) - \bar{B}(t)$ mean-reverts to zero by fraction $\theta^2 \tau$ as a result of each bet. Since bets occur at rate γ per day, the γ from (28) shows that the error $B(t) - \bar{B}(t)$ mean-reverts to zero at rate

$$\rho := \sigma^2 \cdot \Sigma^{-1} \quad (29)$$

per day. Holding volatility constant, resiliency ρ is larger in more efficient markets with smaller $\Sigma^{1/2}$.

Invariance Theorem. *In the meta-model, the number of bets per day γ , their size \tilde{Q} , liquidity L , efficiency $(1/\Sigma)^{1/2}$, and resilience ρ are related to price P , share volume V , volatility σ , and trading activity $W = P \cdot V \cdot \sigma$ by the following invariance relationships, which are consistent with the conjectured invariance relationships in equations (5), (6), (8), (28), and (29):*

$$\gamma = \left(\frac{\lambda \cdot V}{\sigma \cdot P \cdot m} \right)^2 = \left(\frac{E\{\tilde{Q}\}}{V} \right)^{-1} = \frac{(\sigma \cdot L)^2}{m^2} = \frac{\sigma^2}{\theta^2 \tau \cdot \Sigma} = \frac{\rho}{\theta^2 \tau} = \left(\frac{W}{m \cdot \bar{C}_B} \right)^{2/3}. \quad (30)$$

The risk transferred by a bet \tilde{Q} in business time \tilde{I} satisfies the following equation:

$$\tilde{I} := \frac{P \cdot \tilde{Q} \cdot \sigma}{\gamma^{1/2}} = \frac{\tilde{Q}}{V} \cdot W^{2/3} \cdot (m \cdot \bar{C}_B)^{1/3} = \bar{C}_B \cdot \tilde{i}. \quad (31)$$

Here, τ is precision of a signal, θ is fraction of information \tilde{i} incorporated by informed traders, \bar{C}_B is expected cost of a bet, and $m := E\{\tilde{Q}\} / \{E\{\tilde{Q}^2\}\}^{1/2} = E\{\tilde{i}\}$.

Proof of Invariance Relationships. Using equation (17), we write equation (21) for daily volume, equation (24) for expected costs, and Kalman-filtering equation (26) as a system of three equations:

$$\gamma \cdot E\{|\tilde{Q}|\} = V, \quad (32)$$

$$\bar{C}_B = \lambda \cdot E\{\tilde{Q}^2\}, \quad (33)$$

$$\gamma \cdot \lambda^2 \cdot E\{\tilde{Q}^2\} = P^2 \cdot \sigma^2. \quad (34)$$

In the three equations (32), (33), and (34), think of γ , λ , $E\{\tilde{Q}^2\}$, and $E\{|\tilde{Q}|\}$ as endogenous variables and V , \bar{C}_B , P , and σ as exogenous. Since there are three equations and four unknowns, we need a fourth equation. Using a normal distribution for \tilde{Q} , the fourth equation is the moment ratio $m = E\{|\tilde{Q}|\} / \{E\{\tilde{Q}^2\}\}^{1/2}$. Since \tilde{Q} is approximately normally distributed in our meta-model, we have $m \approx (\pi/2)^{1/2}$. For different distributions in different models, m will take different values. If we think of m as an exogenous parameter, we now have four equations in four unknowns.

Using the definition of m and the definition of trading activity $W = P \cdot V \cdot \sigma$, we can solve equations (32), (33), and (34) for γ , $E\{|\tilde{Q}|\}$, and λ , as follows. Multiply the product of (32) and (33) by the square root of (34) and solve for γ to obtain

$$\gamma = \left(\frac{1}{m \cdot \bar{C}_B} \right)^{2/3} \cdot W^{2/3}. \quad (35)$$

Divide the product of (34) and the square of (33) by (32) and solve for $E\{|\tilde{Q}|\}$ to obtain

$$E\{|\tilde{Q}|\} = (m \cdot \bar{C}_B)^{2/3} \cdot V \cdot W^{-2/3}. \quad (36)$$

Divide the product of (34) and the square root of (33) by (32) and solve for λ to obtain

$$\lambda = \left(\frac{m^2}{\bar{C}_B} \right)^{1/3} \cdot \frac{1}{V^2} \cdot W^{4/3}. \quad (37)$$

Equation (36) implies that the measure of illiquidity $1/L$ is

$$\frac{1}{L \cdot \sigma} = \left(\frac{W}{m \cdot \bar{C}_B} \right)^{-1/3} \cdot m^{-1}. \quad (38)$$

Equation (28) and equation (29) imply that market efficiency $1/\Sigma^{1/2}$ and resilience ρ are

$$\rho = \left(\frac{\Sigma^{1/2}}{\sigma} \right)^{-2} = \left(\frac{W}{m \cdot \bar{C}_B} \right)^{2/3} \cdot \frac{1}{\theta^2 \tau}. \quad (39)$$

Define a bet's risk transfer in business time as $\tilde{I} := P\tilde{Q} \cdot \sigma \gamma^{-1/2}$. Equations (15), (17), (28), (35), and (37) imply (31).

The equation (35) for γ , equation (36) for \tilde{Q} , and equation (38) for $1/L$, (39) for $1/\Sigma^{1/2}$ and ρ are summarized in (30). They are respectively equivalent to equations (5), (6), (8), (13), and (14) implied by market microstructure invariance hypothesis in section 1, given that $E\{\tilde{i}\} = m$ implies $E\{\tilde{I}\} = m \cdot \bar{C}_B$ from (31).

Empirically, since trading activity W and its components are observable, we can immediately infer values of γ , $E\{|\tilde{Q}|\}$, $1/L$, λ , $1/\Sigma^{1/2}$ and ρ from equation (30), if the value of constants $m \cdot \bar{C}_B$ and $\theta^2\tau$ are known.

Notes on Existence and Uniqueness of Equilibrium. We define an “equilibrium” of the meta-model as a “steady state” solution to the endogenous variables in terms of the exogenous variables, given an initial price $\bar{B}(t)$.

The roadmap to a more formal proof of existence and uniqueness of an equilibrium to the meta-model is as follows. The meta-model has eight exogenous variables: the share float N , the turnover rate by noise traders η , fundamental volatility σ , Brownian motion $B(t)$ driving fundamental value, the cost of a signal c_I , signal precision τ , the fraction of information incorporated into prices θ , and a moment ratio equal to $m := (\pi/2)^{1/2}$ for a normal distribution of signal \tilde{i} . The meta-model has eight endogenous variables: volume V , trading activity W , the arrival rate of bets γ , average bet size $E|\tilde{Q}|$, market depth λ , market liquidity L , market efficiency $1/\Sigma^{1/2}$, and market resilience ρ . Also, determined endogenously are the bet size distribution of \tilde{Q} and the distribution of the invariant \tilde{I} . Given an initial value of $\bar{B}(t)$, one can calculate an initial price P , $V = \eta N / (1 - \theta)$, $W = \sigma \cdot P \cdot V$, and $\bar{C}_B = \theta c_I / (1 - \theta)$. The equations describing the meta-model then uniquely determine, as described in equations (30), all other endogenous stock-specific variables in terms of W and its components, specifically V , P , and σ , given \bar{C}_B , m , and $\theta^2\tau$.

Note that as P changes, market capitalization changes, and all of the other endogenous variables change. Thus, our “steady state” is really an approximation to something more complicated. At any given time, the market will be moving towards a steady state implied by P since Σ will only change gradually. The market will never reach this steady state because P is constantly changing due to the dynamics implied by the model. In particular, market velocity γ is changing randomly, leading to something analogous to time deformation. If velocity increases, returns volatility may temporarily increase, as increased pricing accuracy makes the market more efficient. In the long run, however, returns volatility is pinned down by fundamental volatility, and changes in velocity show up in a persistent way as changes in market liquidity. In the limit as market capitalization becomes large, we conjecture (but leave a formal proof for future research) that the market becomes so efficient that it is “very close” to a steady state.

Invariance of Bets. In the definition of \tilde{I} in equation (31), P is measured in dollars per share, \tilde{Q} is measured in shares, σ is measured per square root of time, and γ is measured per time. This implies \tilde{I} is measured in dollars, the same units as \bar{C}_B . In

equation (31) therefore \tilde{i} stands for a unitless distribution of information.

The basic microstructure invariance hypothesis states that the distribution of \tilde{I} does not vary across stocks or across time. The meta-model reveals that microstructure invariance is ultimately connected to granularity of information structure and invariance of expected costs and profits per signal across markets. Signals \tilde{i} are drawn from distribution with zero mean and variance of one. If the shape of information distribution is the same across markets (not precision of signals!), then the basic invariance hypothesis (31) is equivalent to the hypothesis that \bar{C}_B does not vary across stocks or time.

Although invariance relationships depend on the assumption that the cost of executing a bet \bar{C}_B is constant across stocks, \bar{C}_B is not the “deepest” parameter in this model. Since equation (18) and equation (24) imply that $\bar{\pi}_I = \theta \bar{C}_B$, we plug $\bar{\pi}_I$ into equation (25) and find $\bar{C}_B = \theta c_I / (1 - \theta)$. The value of \bar{C}_B is constant across stocks if c_I and θ are constant across stocks. It is useful to think of the cost of private information c_I as proportional to the average wages of finance professionals, adjusted for their productivity. The productivity-adjusted wage of a finance professional is therefore a “deeper” parameter than the endogenous cost of executing a bet \bar{C}_B . Ultimately, invariance is based on the idea that the effort c_I required to generate one discrete bet has a distribution which does not vary across stocks (even if the precision $\theta^2 \cdot \tau$ of the resulting information revealed in prices does vary across stocks). The invariance of information costs c_I leads to granularity of information flow, which is embedded in invariance relationships. Why this might be so takes us beyond the scope of this paper.

Invariance of Transactions Costs. The meta-model does not require any assumptions about the functional form of trading costs $C_B(\tilde{Q})$ as a function of bet size \tilde{Q} . The meta-model only requires restrictions on the expected transactions cost of a typical bet $\bar{C}_B = E\{C_B(\tilde{Q})\}$ in equation (24). The meta-model is consistent with many functional forms of $C_B(\cdot)$ implied by invariance.

Suppose, for example, that $C_B(\tilde{Q}) = a \cdot |\tilde{Q}|^{r+1} + b \cdot |\tilde{Q}|$; then the average transactions cost function $C(\tilde{Q}) = C_B(\tilde{Q}) / (P|\tilde{Q}|)$ is

$$C(\tilde{Q}) = \frac{a \cdot |\tilde{Q}|^{r+1} + b \cdot |\tilde{Q}|}{P|\tilde{Q}|}. \quad (40)$$

Suppose also there is a constant α such that $a \cdot E\{|\tilde{Q}|^{r+1}\} = \alpha \cdot \bar{C}_B$ and $b \cdot E\{|\tilde{Q}|\} = (1 - \alpha) \cdot \bar{C}_B$. Using invariance relation (31), we have

$$C(\tilde{Q}) = \sigma W^{-1/3} \cdot \bar{\kappa}_0 + \sigma W^{(2r-1)/3} \cdot \left| \frac{\tilde{Q}}{\bar{V}} \right|^r \cdot \bar{\kappa}_I, \quad (41)$$

where $\bar{\kappa}_0 := (m\bar{C}_B)^{1/3}/m$ and $\bar{\kappa}_I := (m\bar{C}_B)^{(r+1)/3} \bar{C}_B / E\{|\tilde{I}|^{r+1}\}$. If $r = 1$, then equation (41) is equivalent to the linear price impact function (11) implied by the

market microstructure invariance hypothesis in section 1. If $r = 1/2$, then equation (41) is equivalent to the square root impact function (12) in section 1.

Neither the invariance hypothesis nor the meta-model provide intuition on the functional form of the transactions cost function. In section 6, we approach this question empirically and calibrate its form using portfolio transition data. We find that a square root function ($r = 1/2$) with a constant bid-ask spread term better describes the data than linear function ($r = 1$) with a constant bid-ask spread term.

Note that the cost functions $C_B(\tilde{Q})$ does not specify functional forms of its permanent and temporary components. The meta-model implies that linear long-term permanent impact λ is such that, on average, exactly a half of $E\{C_B(\tilde{Q})\}$ is attributed to permanent impact costs $\lambda E\tilde{Q}^2/2$. If we wish to write $C_B(Q) = C_P(Q) + C_T(Q)$, then the meta-model implies $C_P(Q) = \lambda\tilde{Q}^2/2$ and $C_T(Q) = C_B(Q) - \lambda\tilde{Q}^2/2$, with $E\{C_B(\tilde{Q})\}/2 = \lambda E\tilde{Q}^2/2$.

Invariance of Market Efficiency and Resilience. Inferring our measures of market efficiency $(1/\Sigma)^{1/2}$ and resilience ρ from equation (30) requires not only observing W and knowing $m \cdot \bar{C}_B$ but also knowing $\theta^2 \cdot \tau$, which measures the information content of a bet when it is not know whether the bet is informed or noise. If $\theta^2 \cdot \tau$ is invariant across stocks and time, then $\sigma/\Sigma^{1/2}$ is proportional to $\gamma^{1/2}$ and ρ is proportional to γ . The bigger is γ , the more efficient and resilient is the market. Whether or not $\theta^2 \cdot \tau$ is invariant across stocks is an interesting area for future research.

Recall that Black (1986) expressed the idea that “an efficient market is one in which price is within a factor 2 of value.” If deviation from the factor of 2 is interpreted as a $\xi = 1.0$ standard deviation event, then his statement is consistent with the interpretation $1.0 \cdot \Sigma^{1/2} = \log_e 2 \approx 0.70$. If a stock’s annual volatility is about 35%, then Fischer Black’s measure of market efficiency further implies that prices are about $(0.70/0.35)^2 = 4$ years “behind” fundamental value. We conjecture that future empirical research may indicate that financial markets are much more efficient than conjectured by Black.

Since efficiency and resilience are closely related, it might be easier empirically to calibrate Black’s measure of market efficiency by examining how fast “permanent” price effects from noise trades die out over time rather than measuring directly how accurately prices approximate fundamental value. The error $B(t) - \bar{B}(t)$ mean-reverts to zero at rate $\rho = \sigma^2 \cdot \Sigma^{-1}$ equal to $(0.35/0.70)^2 = 0.25$ per year. This implies the half-life of the “permanent” price impact of a noise trade is $\log_e 2/\rho$, i.e., half of the price impact of a noise trade dies out in about $\log_e 2/0.25 \approx 2.8$ years.

The meta-model adds additional structure that imposes restrictions on the three invariance hypotheses in section 1. The meta-model implies that $E|\tilde{I}|/m = \bar{C}_B$, see equation (31), or, equivalently, that the standard deviation of \tilde{I} is equal to \bar{C}_B . Intuitively, this assumption is consistent with the intuition that market makers must break even trading against informed bets which have permanent price impact. This additional assumption imposes a particular structure on the proportionality constants

in invariance relationships (5), (6), (9), (13), and (14) and allows us to write those disconnected relationships in a consolidated form of the invariance theorem.

3 Microstructure Invariance in the Context of Market Microstructure Literature

Microstructure invariance does not undermine or contradict other theoretical models of market microstructure. Instead, it builds a bridge from theoretical models to empirical tests of those models. Theoretical microstructure models, including the meta-model in this paper, use the idea that order flow imbalances move prices to construct measures of market depth or liquidity. Microstructure invariance imposes cross-sectional restrictions which make it easier to implement liquidity measures based on order flow imbalances.

Many theoretical models use game theory to model trading. These models typically make specific assumptions about the consistency of beliefs across traders, the flow of public and private information which informed traders use to trade, the flow of orders from liquidity traders, and auction mechanisms in the context of which market makers compete to take the other side of trades. Some models emphasize adverse selection, such as Treynor (1995), Kyle (1985), Glosten and Milgrom (1985), and Back and Baruch (2004); some models emphasize inventory dynamics, such as Grossman and Miller (1988) and Campbell and Kyle (1993); some models emphasize both, such as Grossman and Stiglitz (1980) and Wang (1993).

Theoretical models suggest that order flow imbalances move prices. Parameters describing how order flow moves prices depend on the specifics of each model. Understanding the cross-sectional empirical implications concerning how order flow moves price has been difficult. Theoretical microstructure models provide neither a unified framework for mapping the theoretical concept of an order flow imbalance into its empirical measurements nor precise predictions concerning how price impact varies across different stocks.

Instead, researchers take a purely empirical approach. They regress price changes on imperfect empirical proxies for order flow imbalances—e.g., the difference between uptick and downtick volume, popularized by Lee and Ready (1991)—to obtain market impact coefficients, which they relate to stock characteristics such as market capitalization, trading volume, and volatility. Breen, Hodrick and Korajczyk (2002) is an example of this approach. There is a voluminous empirical literature describing how the rate at which orders arrive in calendar time, the dollar size of orders, the market impact costs, and bid-ask spread costs vary across different assets. For example, Brennan and Subrahmanyam (1998) estimate order size as a function of various stock characteristics. Hasbrouck (2007) is a good survey of empirical literature.

By contrast, microstructure invariance generates precise and empirically testable predictions about how the size of bets, arrival rate of bets, market impact costs, and

bid-ask spread costs vary across assets with different levels of trading activity. These predictions are based on a common intuition shared by many models. The usually unidentifiable parameters of theoretical models show up as invariant constants (e.g., $E\{|\tilde{I}|\}$ and \tilde{C}_B), which can be calibrated from data. Microstructure invariance is a modeling principle applicable to different models, not a model itself. It compliments theoretical models by making it easier to test them empirically.

Example: Model of Kyle (1985). Consider, for example, the continuous time theoretical model of Kyle (1985). The market depth formula $\lambda = \sigma_V/\sigma_U$ in that model measures market depth (in units of dollars per share per share) as the ratio of the standard deviation of stock price changes σ_V (measured in dollars per share per unit of time) to the standard deviation in order flow imbalances σ_U (measured in shares per unit of time). This formula asserts that price fluctuations result from the linear impact of order flow imbalances. It does not depend on specific assumptions about interactions among market makers, informed traders and noise traders. An empirical implementation of the market impact formula $\lambda = \sigma_V/\sigma_U$ should *not* be considered a test of the specific assumptions of the model of Kyle (1985), such as the existence of a monopolistic informed trader who trades smoothly and patiently in a context where less patient liquidity traders trade more aggressively and market makers set stock prices efficiently. Instead, empirical implementation of the formula $\lambda = \sigma_V/\sigma_U$ attempts the more general task of measuring a market impact coefficient λ based on the assumption that price fluctuation result from the linear impact of order flow innovations, a property shared by many models.

Measuring the numerator σ_V is much more straightforward than measuring the denominator σ_U . The value of σ_V is easily inferred from a stock price and returns volatility, under the maintained hypothesis that risk-neutral market makers make markets semi-strong form efficient. In the context of our meta-model, $\sigma_V = \bar{\sigma} \cdot P$.

Measuring the denominator σ_U is difficult because the connection between observed trading volume and order flow imbalances is not straightforward. Intuitively, σ_U should be related to trading volume in some way. The continuous-time model provides no help concerning what this relationship is. In the Brownian motion model of Kyle (1985), trading volume is infinite. Without some approach for measuring σ_U , the model is untestable. In the context of our meta-model, order flow imbalances result from random discrete decisions by traders to change stock holdings and a standard deviation of order imbalances is equal to $\sigma_U = \gamma^{1/2} \cdot (E\tilde{Q}^2)^{1/2}$. This calculation is also consistent with the spirit of other models, such as Glosten and Milgrom (1985).

The formulas for the numerator and denominator imply that the price impact of a bet of \tilde{Q} shares, expressed as a fraction of the value of a share, is given by

$$\frac{\lambda \cdot X}{P} = \frac{\sigma_V}{\sigma_U} \cdot \frac{X}{P} = \bar{\sigma}\gamma^{-1/2} \cdot \frac{X}{(E\tilde{Q}^2)^{1/2}}. \quad (42)$$

Thus, a one-standard deviation bet event has a price impact $\bar{\sigma}\gamma^{-1/2}$ equal to one stan-

dard deviation of returns volatility measured over a time interval $1/\gamma$ corresponding to the expected time between bet arrivals.

This formula can be tested empirically using restrictions imposed by microstructure invariance. Using equations (5) and (6) to determine how γ and moments of \tilde{Q} vary with observable volume and volatility, we find

$$\lambda := \frac{\bar{\sigma}}{\gamma^{-1/2} \cdot (E\tilde{Q}^2)^{1/2}} = \frac{\bar{\sigma}}{\bar{V}} \cdot \bar{W}^{1/3} \cdot [E\{|\tilde{I}|^2\}]^{-1/2} \cdot \iota^2. \quad (43)$$

The final step is to calibrate a constant $[E\{|\tilde{I}|^2\}]^{-1/2} \cdot \iota^2$, which does not depend on units of time, because \tilde{I} is measured in dollars.

As an alternative to invariance, the formula $\lambda = \sigma_V/\sigma_U$ can be implemented empirically by imposing different assumptions concerning the connection between σ_U and trading volume. For example, we can think of the illiquidity ratio in Amihud (2002) as an empirical implementation of that formula. Amihud’s illiquidity ratio is the time-series average of the daily ratios of absolute value of returns to dollar volume. To the extent that dollar volume is relatively stable across time and returns are drawn from the same distribution, this measure is effectively proportional to $\sigma/(PV)$. This would imply that Amihud’s implementation is effectively based on the assumption that σ_U is proportional to volume V , as if the expected arrival rate of orders in Kyle (1985) were some unknown constant, the same for all stocks. The ratio $\sigma/(PV)$ has time units—unlike our illiquidity measure $1/L = \iota^2 \bar{C}_B \cdot [P\bar{V}/\bar{\sigma}^2]^{-1/3}$. The proportionality constant in Amihud’s illiquidity ratio depends therefore on the units in which time is measured. Invariance implies that Kyle (1985) should be implemented as a model about much longer periods of time for less active stocks, implying that the proportionality constants might have to be different across stocks.

4 Data

Portfolio Transitions Data. We test the empirical implications of market microstructure invariance using a proprietary dataset of portfolio transitions from a leading vendor of portfolio transition services. During the evaluation period, this portfolio transition vendor supervised more than 30 percent of outsourced U.S. portfolio transitions. The sample includes 2,552 portfolio transitions executed over the period 2001-2005 for U.S. clients. A portfolio transition may involve orders for hundreds of individual stocks. Each order is a stock-transition pair potentially executed over multiple days using a combination of internal crosses, external crosses, and open-market transactions.

The portfolio transitions dataset contains fields identifying the portfolio transition; its starting and ending dates; the stock traded; the trade date; the number of shares traded; a buy or sell indicator; the average execution price; the pre-transition benchmark price (closing price the day before the transition trades began); commis-

sions; SEC fees; and a trading venue indicator distinguishing among internal crossing networks, external crossing networks, open market transactions, and in-kind transfers.

When old “legacy” and new “target” portfolios overlap, positions are transferred from the legacy to the new portfolio as “in-kind” transfers. For example, if the legacy portfolio holds 10,000 shares of IBM stock and the new portfolio holds 4,000 shares of IBM, then 4,000 shares are transferred in-kind and the balance of 6,000 shares is sold. The in-kind transfers do not incur transactions costs and have no effect on our empirical analysis.

We augment the portfolio transitions data with stock price, returns, and volume data from CRSP. Only common stocks (CRSP share codes of 10 and 11) listed on the New York Stock Exchange (NYSE), the American Stock Exchange (Amex), and NASDAQ in the period of January 2001 through December 2005 are included in the sample. ADRs, REITS, and closed-end funds are excluded. Also excluded are stocks with missing CRSP information necessary to construct variables used for empirical tests, transition orders in high-priced Berkshire Hathaway class A shares, and transition observations which appeared to contain typographical errors and obvious inaccuracies. Since it is unclear from the data whether adjustments for dividends and stock splits are made in a consistent manner across all transitions, all observations with non-zero payouts during the first week following the starting date of portfolio transitions were excluded from statistical tests.

After exclusions, there are 439,765 observations (“orders”), including 201,401 buy orders and 238,364 sell orders.

CRSP Data: Prices, Volume, and Volatility. For each of the transition-stock observations ($i = 1, \dots, 439765$), we collect data on pre-transition benchmark price, expected volume, and expected volatility. The benchmark price, denoted $P_{0,i}$, is the closing price for the stock the evening before the first trade is made in any of the stocks in the portfolio transition. A proxy for expected daily trading volume, denoted V_i (in shares), is the average daily trading volume for the stock in the previous full pre-transition calendar month.

The expected volatility of daily returns, denoted σ_i for order i , is calculated using past daily returns in two different ways.

First, for each security j and each calendar month m , we estimate the monthly standard deviation of returns $\sigma_{j,m}$ as the square root of the sum of squared daily returns for the full calendar month m (without de-meaning or adjusting for autocorrelation). We define $\sigma_i = \sigma_{j,m}/N_m^{1/2}$, where j corresponds to the stock traded in order i , m is the previous full calendar month preceding order i , and N_m is the number of CRSP trading days in month m .

Second, to reduce effects from the positive skewness of the standard deviation estimates, we estimate for each stock j a third-order moving average process for the changes in $\ln(\sigma_{j,m})$ for all months m over the entire period 2001-2005: $(1 - L)\ln(\sigma_{j,m}) = \Theta_{j,0} + (1 - \Theta_{j,1}L - \Theta_{j,2}L^2 - \Theta_{j,3}L^3)u_{j,m}$. Letting $y_{j,m}$ denote the estimate

of $\ln(\sigma_{j,m})$ and \hat{V}_j the variance of the prediction error, we alternatively define the conditional forecast for the volatility of daily returns by $\sigma_i = \exp(y_{j,m} + \hat{V}_j/2)/N_m^{1/2}$, where m is the *current* full calendar month for order i .

As noisy estimates of true volatility, these proxies may potentially introduce an error-in-variables problem into the regressions below. While we report below results using the second definition of σ_i based on the log-ARIMA model, these results remain quantitatively similar when we use the first definition of σ_i based on simple historical volatility during the preceding full calendar month.

Except to the extent that the ARIMA model uses in-sample data to estimate model parameters, we use the pre-transition variables known to the market before portfolio transition trades are executed in order to avoid any spurious effects from using contemporaneous variables.

Descriptive Statistics. Table 1 reports descriptive statistics for traded securities in panel A and for individual transition orders in panel B. The first column reports statistics for all securities in aggregate; the remaining ten columns report statistics for stocks in ten dollar volume groups. Instead of dividing the securities into ten deciles with the same number of securities in each decile, volume break points are set at the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of trading volume for the universe of stocks listed on the NYSE with CRSP share codes of 10 and 11. Group 1 contains stocks in the bottom 30th percentile by dollar trading volume. Group 10 approximately corresponds to the universe of S&P 100 stocks. The top five groups approximately cover the universe of S&P 500 stocks. Smaller percentiles for the more active stocks make it possible to focus on the stocks which are most important economically. Each month the thresholds are recalculated and the stocks are reshuffled across bins.

Panel A of table 1 reports descriptive statistics for traded securities. For the entire sample, the median daily volume is \$18.72 million, ranging from \$1.13 million for the lowest volume group to \$212.85 million for the highest volume group. The median volatility is 1.93 percent per day, ranging from 1.76 percent in the highest volume decile to 2.16 in the lowest decile. Since there is so much more cross-sectional variation in dollar volume than in volatility across stocks, the variation in trading activity across stocks is related mostly to variation in dollar volume. Trading activity differs by a factor of 150 between stocks in the lowest group and stocks in the highest group, and this variation creates statistical power helpful in determining how transactions costs and order sizes vary with trading activity.

The median quoted bid-ask spread, obtained from the transition dataset, is 12.04 basis points; its mean is 25.42 basis points. From lowest volume group to highest volume group, the median spread declines monotonically from 40.96 to 4.83 basis points, by a factor of 8.48. A back-of-the-envelope calculation based on invariance suggests that spreads should decrease approximately by a factor of $150^{1/3} \approx 5.31$ from lowest to highest volume group. The difference between 5.31 and 8.48 warrants

further investigation. The monotonic decline of almost one order of magnitude is potentially large enough to generate significant statistical power in estimates of a bid-ask spread component of transactions costs based on implementation shortfall.

Panel B of table 1 reports properties of portfolio transition order sizes. The average order size is 4.20% of average daily volume, declining monotonically across the ten volume groups from 16.23% in the smallest group to 0.49% in the largest group, by a factor of 33.12. The median order is 0.57% of average daily volume, also declining monotonically from 3.33% in the smallest group to 0.14% in the largest group, by a factor of 23.79. The invariance hypothesis implies that order sizes should decline by a factor of approximately $150^{2/3} \approx 28.23$, a value which matches the data closely. The medians are much smaller than the means, indicating that distributions of order sizes are skewed to the right. We show below that the distribution of order sizes fits closely a log-normal.

The average trading cost (based on implementation shortfall as explained below) is 16.79 basis points per order, ranging from 44.95 basis points in lowest volume group to 6.16 basis points in the highest group (excluding commissions and SEC fees) by a factor of 7.30. Invariance implies that costs should fall by a factor of $150^{1/3} \approx 5.31$, somewhat smaller than the actual decline. Note that the SEC fee represents a cost of about 0.29 basis points and does not vary much across volume groups. The average commission is 7.43 basis points, declining monotonically from 14.90 basis points for the lowest group to 2.68 basis points for the highest group. Since commission may be negotiated for the entire transition, the allocation of commission costs to individual stocks is an accounting exercise with little economic meaning.

Portfolio transitions consist of orders for dozens or hundreds of stocks, which are typically executed over several days. About 60% of orders are executed during the first day of the portfolio transition. Since transition managers often operate under a “cash-in-advance” constraint—using proceeds from selling stocks in a legacy portfolio to acquire stocks in a target portfolio—sell orders tend to be executed slightly faster than buy orders (1.72 days versus 1.85 days). In terms of dollar volume, about 41%, 23%, 15%, 7% and 5% of dollar volume is executed on the first day through the fifth days, respectively. The two longest transitions in the sample were executed over 18 and 19 business days. The time frame for a portfolio transition is usually set before its actual implementation begins.

5 Empirical Tests Based on Order Sizes

Market microstructure invariance predicts that the distribution of $\bar{W}^{2/3} \cdot \tilde{Q}/\bar{V}$ does not vary across stocks or time (see equation (6)). We test these predictions using data on portfolio transition orders, making the identifying assumption that portfolio transition orders are proportional to bets.

Portfolio Transitions and Bets. Since bets are statistically independent intended orders, bets can be conceptually difficult for researchers to observe. Consider, for example, a trader who makes a decision on Monday to make one bet to buy 100,000 shares of stock, then implements the bet by purchasing 20,000 shares on Monday and 80,000 shares on Thursday. To an econometrician, this one bet for 100,000 shares may be difficult to distinguish from two bets for 20,000 shares and 80,000 shares, respectively. Portfolio transitions make the task of identifying bets easier because the size of each order in this example is known and recorded on Monday, even if the order is executed over several subsequent days.

Portfolio transition orders may not have a size distribution matching precisely the size distribution of typical bets. Transition orders may be smaller than bets if transitions tend to liquidate a portion of an asset manager’s positions or larger than bets if transitions liquidate the sum of bets made by the asset manager in the past. When both target and legacy portfolios hold long positions in the same stock, the portfolio transition order may represent the difference between two bets.

Let X_i denote the unsigned number of shares transacted in portfolio transition order i , $i = 1, \dots, 439765$. The quantity X_i sums shares traded over multiple days, excluding in-kind transfers.

We make the identifying assumption that, for some constant δ which does not vary across stocks with different characteristics such as trading activity W , the distribution of scaled portfolio transition orders $\delta \cdot X_i$ is the same as the distribution of a typical bet in the same stock at the same time, denoted $|\tilde{Q}_i|$.

Invariant Order Size Distribution as Log-Normal. Let $W_i := V_i \cdot P_i \cdot \sigma_i$ and $\bar{W}_i := \bar{V}_i \cdot P_i \cdot \bar{\sigma}_i$ denote trading activity and bet activity for the stock in transition order i , respectively. Under the identifying assumption that portfolio transition orders are proportional to bets, invariance of the theoretical distribution of bets adjusted for bet activity $\bar{W}^{2/3} \cdot |\tilde{Q}|/\bar{V}$ from equation (6) implies that the empirical distribution of $\bar{W}_i^{2/3} \cdot \delta X_i/\bar{V}_i$ does not vary with stock characteristics such as volume, volatility, stock price, or market capitalization.

To facilitate intuitive interpretation of parameter estimates, we scale observations by a hypothetical “benchmark stock” with share price P^* of \$40, daily volume V^* of one million shares, and volatility σ^* of 2% per day, implying $W^* = 40 \cdot 10^6 \cdot 0.02$. Table 1 implies that this benchmark stock would belong to the bottom tercile of S&P 500, i.e., to volume group 7.

Combining the identifying assumption with equations (1) with (2) to convert the bet activity variables \bar{W}_i and \bar{V}_i to trading activity variables W_i and V_i and taking logs, invariance implies

$$\ln \left(\left[\frac{W_i}{W^*} \right]^{2/3} \cdot \frac{X_i}{V_i} \right) = \ln(\bar{q}) + \tilde{\epsilon}, \quad (44)$$

where $\ln(\bar{q}) := E\{\ln(|\tilde{Q}^*|/V^*)\} - 1/3 \ln(\zeta/\zeta^*) - 2/3 \ln(\psi/\psi^*) - \ln(\delta)$ and $\tilde{\epsilon} := \ln(|\tilde{I}|) -$

$E\{\ln(|\tilde{I}|)\}$ has a zero-mean invariant distribution. Under the identifying assumptions that the volume multiplier ζ , the volatility multiplier ψ and the portfolio transition size multiplier δ do not vary across observations, $\ln(\bar{q})$ is an invariant constant $\ln(\bar{q}) = E\{\ln(|\tilde{Q}^*|/V^*)\} - \ln(\delta)$. Multiplication of X_i/V_i by $(W_i/W^*)^{2/3}$ scales each observation on the left-side so that it has the same invariant distribution as the log of a hypothetical portfolio transition order in the benchmark stock, expressed as a fraction of its expected daily volume. If $\delta = 1$, the distribution of portfolio transition orders matches the distribution of bets.

As we examine this hypothesis, we will also examine the stronger log-normality hypothesis—not implied by microstructure invariance—that $\tilde{\epsilon}$ has a normal distribution, which implies invariant log-normal distribution for order sizes adjusted for trading activity $W^{2/3} \cdot |\tilde{Q}|/V$.

Under the assumption of log-normality, the predictions of equation (44) can be calibrated by two parameter estimates, the sample of mean and the sample variance. The pooled sample mean of -5.71 is an estimate of $\ln(\bar{q})$, and the pooled sample variance of 2.53 is an estimate of the variance of $\tilde{\epsilon}$. Since $\exp(-5.71) \approx 0.0033$, log-normality implies a median portfolio transition order size of 0.33% of expected daily volume for the benchmark stock. Since $\exp(2.53^{1/2}) \approx 4.90$, a one standard deviation increase in order size is a factor of 4.90 for all stocks.

There is not much difference in the distributions of buy and sell orders. For buy orders, the mean is -5.70 for the variance is 2.51 ; for sell orders, the mean is -5.71 and the variance is 2.55 .

Order Size Distribution for Volume and Volatility Groups. Log-normality and invariance of bets imply that the distribution of $2/3 \cdot \log(W_i/W^*) + \log(X_i/V_i)$ should remain approximately $N(-5.71, 2.54)$ for subsets of orders sorted into different groups based on characteristics such as dollar volume, volatility, stock price, and turnover.

To examine this hypothesis visually, we plot the empirical distributions of the left-hand side of equation (44), $2/3 \cdot \ln(W_i/W^*) + \ln(X_i/V_i)$, for selected volume and volatility groups. As before, we define ten dollar volume groups with thresholds corresponding to the *30th*, *50th*, *60th*, *70th*, *75th*, *80th*, *85th*, *90th*, and *95th* percentiles of NYSE dollar volume. We define five volatility groups with thresholds corresponding to the *20th*, *40th*, *60th*, and *80th* percentiles of returns standard deviation for NYSE stocks. On each plot, we superimpose the bell-shaped density function $N(-5.71, 2.53)$ matching the mean and variance of the pooled sample.

Figure 2 shows plots of the empirical distributions of $\ln[X_i/V_i] + 2/3 \cdot \ln[W_i/W^*]$ for volume groups 1, 4, 7, 9, and 10 and for volatility groups 1, 3 and 5. Consistent with the invariance hypothesis, these fifteen distributions of W -adjusted order sizes are all visually strikingly similar to the normal distribution with the pooled mean and variance $N(-5.71, 2.53)$. Results for the remaining 35 subgroups also look very similar and therefore are not presented in this paper. The visual similarity of the

distributions is reflected in the similarity of their first four moments. For the 15 volume-volatility groups, the means range from -6.03 to -5.41 , close to the mean of -5.71 for the pooled sample. The variances range from 2.23 to 2.90 , also close to the variance of 2.53 for the pooled sample. The skewness ranges from -0.21 to 0.10 , close to skewness of zero for the normal distribution. The kurtosis ranges from 2.73 to 3.38 , also close to the kurtosis of 3 for a normal random variable. These results suggest that the assumption that order sizes are distributed as log-normal random variables is a reasonable one. Scaling order sizes by $(W/W^*)^{2/3}$, as implied by the invariance hypothesis, adjusts the means of the distributions so that they do appear visually to be similar.

Despite the visual similarity, a Kolmogorov-Smirnov test rejects the hypothesis that all fifty empirical distributions are generated from the same normal distribution. The standard deviation of the means across bins is larger than implied by a common normal distribution. Microstructure invariance does not describe the data perfectly, but it is close enough to suggest that invariance makes a good benchmark from which the modest deviations seen in these plots can be investigated in future research.

Figure 3 further examines log-normality by focusing on the tails of the distributions of portfolio transition orders. For each of the five volume groups 1, 4, 7, 9, and 10, panel A shows quantile-quantile plots of the empirical distribution of $\ln[X_i/V_i] + 2/3 \cdot \ln[W_i/W^*]$ versus a normal distribution with the same mean and variance. The more similar these empirical distributions are to a normal distribution, the closer the plots should be to the 45-degree line. Panel B shows logs of ranks based on scaled order sizes. Under the hypothesis of log-normality, the right tail should be quadratic. A straight line in the right tail implies a power law. Both panels show that the empirical distributions are similar to a normal distribution, except in the far right and left tails.

In panel A, the smallest orders in the left tails tend to be smaller than implied by a normal distribution. These observations are economically insignificant. Most of them represent one-share transactions in low-price stocks (perhaps the result of coding errors in the data). There are too few such orders to have a meaningful effect on our statistical results.

In panel A, the largest orders in the right tails are much more important economically. On each subplot, a handful of positive outliers (out of 400,000+ observations) do not appear to fit a normal distribution. The largest orders in low-volume stocks appear to be smaller than implied by a normal distribution, and the largest orders in high-volume stocks appear to be larger than implied by a normal distribution.

The finding that the largest orders in low-volume stocks are smaller than implied by a log-normal may be explained by reporting requirements. According to Section 13(d) of the 1934 Act and Regulation 13D, institutions are required to report their holdings whenever they acquire ownership of more than 5 percent of the outstanding shares of publicly traded companies. To avoid reporting requirements, large institutional investors may intentionally acquire fewer shares when intended holdings would

otherwise exceed the 5% reporting threshold. Indeed, all 400,000+ portfolio transition orders are for amounts smaller than 4.5% of shares outstanding. A closer examination reveals that the five largest orders for low-volume stocks accounts for about 2%, 3%, 4%, 4%, and 4% of shares outstanding, respectively, just below the 5% threshold. The largest order in high-volume stocks is for about 1% of shares outstanding.

To summarize, we conclude that the distribution of portfolio transition order sizes appears to conform closely—but not exactly—to the hypothesis of bet size invariance. Furthermore, the distribution of order sizes appears to be quite similar to—but not exactly equal to—a log-normal.

OLS Estimates of Order Size. The order size predictions from equation (44) can be tested using a simple log-linear OLS regression

$$\ln \left[\frac{X_i}{V_i} \right] = \ln [\bar{q}] + \alpha_0 \cdot \ln \left[\frac{W_i}{W^*} \right] + \tilde{\epsilon}. \quad (45)$$

Invariance of bets implies $\alpha_0 = -2/3$. To adjust standard errors of OLS estimates of α_0 for positive contemporaneous correlation in transition order sizes across different stocks, the 439,765 observations are pooled by week over the 2001-2005 period into 4,389 clusters across 17 industry categories. The double clustering by weeks and industries conservatively adjusts standard errors for large portfolio transitions that may involve hundreds of relatively large orders, executed during the course of a week and potentially concentrated in particular industries.

A potential econometric difficulty with the log-linear specification in equation (45) is that taking the log of order size as a fraction of average daily volume may create large negative outliers from tiny, economically meaningless orders, with the inordinately large influence on reported results. Since we have shown above that the shape of the distribution of scaled order sizes closely matches a log-normal, these tiny orders are expected to have only a negligible distorting effect on estimates.

Table 2 presents estimates for the OLS coefficients in equation (45). The first column of the table reports the results of a regression pooling all the data. The four other columns in the table report results for four separate OLS regressions in which the parameters are estimated separately for NYSE Buys, NYSE Sells, NASDAQ Buys, and NASDAQ Sells.

For the entire sample, the estimate for α_0 is $\hat{\alpha}_0 = -0.62$ with standard error of 0.009. Economically, the point estimate for α_0 is close to the value $-2/3$ predicted by the invariance hypothesis, but the hypothesis $\alpha_0 = -2/3$ is strongly rejected ($F = 25.31, p < 0.0001$) because the standard error is very small.

When the sample is broken down into NYSE Buys, NYSE Sells, NASDAQ Buys, and NASDAQ Sells, it is interesting to note that the estimated coefficients for buy orders, -0.63 for NYSE and -0.71 for NASDAQ, are closer to $-2/3$ than the coefficients for sell orders, -0.59 for both NYSE and NASDAQ. Since portfolio transitions tend to be applied to long-only portfolios, sell orders might represent liquidations of

past bets. If the size distribution of sell orders depends on past values of volume and volatility—not current values—there is an errors-in-variables problem which may bias coefficient estimates.

Quantile Estimates of Order Sizes. Table 7 in Appendix presents quantile regression results for equation (45) based on the *1st* (smallest orders), *5th*, *25th*, *50th*, *75th*, *95th*, and *99th* percentiles (largest orders). The corresponding quantile estimates for α_0 are -0.65 , -0.64 , -0.61 , -0.62 , -0.61 , -0.64 , and -0.63 , respectively. Although the hypothesis $\alpha_0 = -2/3$ is rejected due to small standard errors, all quantile estimates are economically close to the value of $-2/3$ predicted by the invariance hypothesis.

Model Calibration and Its Economic Interpretation. To calibrate the hypothesis of bet size invariance with log-normality, we impose the restriction $\alpha_0 = -2/3$ on equation (45). Thus, only the constant term in the regression needs to be estimated. The results of this calibration exercise are presented in table 3. The estimated constant term, -5.71 , is the previously reported sample mean of $\ln(\bar{q})$ in equation (44). The mean-square error, 2.53 is the previously reported sample variance of $\tilde{\epsilon}$ in equation (44).

The log of trading activity $\ln(W/W^*)$, with the coefficient $\alpha_0 = -2/3$ imposed by invariance, explains a significance percentage of the variation of order size as a fraction of volume X_i/V_i ; the R^2 (with zero degrees of freedom) is 0.3149 .

When the parameter α_0 is estimated rather than held fixed, as reported in table 2, changing α_0 from the predicted value of $\alpha_0 = -2/3$ to the estimated value of $\hat{\alpha}_0 = -0.62$ increases the R^2 from $R^2 = 0.3149$ reported in table 3 to $R^2 = 0.3167$ reported in table 2, a modest increase of 0.0018 . Although statistically significant, the addition of one degree of freedom does not add much explanatory power.

We relax the specification further by allowing the coefficients on the three components of trading activity—volatility σ_i , price $P_{0,i}$, and volume V_i —as well as monthly turnover rate ν_i to vary freely:

$$\ln \left[\frac{X_i}{V_i} \right] = \ln [\bar{q}] + \alpha_0 \cdot \ln \left[\frac{W_i}{W^*} \right] + b_1 \cdot \ln \left[\frac{\sigma_i}{0.02} \right] + b_2 \cdot \ln \left[\frac{P_{0,i}}{40} \right] + b_3 \cdot \ln \left[\frac{V_i}{10^6} \right] + b_4 \cdot \ln \left[\frac{\nu_i}{1/12} \right] + \tilde{\epsilon}.$$

This regression imposes on $\ln(W_i/W^*)$ the coefficient $\alpha_0 = -2/3$ predicted by invariance and then allows the coefficient b_1 , b_2 , b_3 , b_4 on the three components of W_i and turnover rate to vary freely. The invariance hypothesis implies $b_1 = b_2 = b_3 = b_4 = 0$. Table 3 reports that increasing the degrees of freedom from one to four increases the R^2 of the regression from $R^2 = 0.3167$ to $R^2 = 0.3229$, an increase of 0.0062 . Although statistically significant, the improvement in R^2 is again modest.

Invariance explains much—but not quite all—of the variation in portfolio transition order size across stocks.

The point estimates for the coefficient on volatility of $\hat{b}_1 = 0.42$, the coefficient on price of $\hat{b}_2 = 0.24$, the coefficient on share volume of $\hat{b}_3 = 0.06$, the coefficient on turnover rate of $\hat{b}_4 = -0.18$ are all statistically significant, with standard errors of 0.040, 0.019, 0.010 and 0.015, respectively (see table 8 in the Appendix). The coefficients on volatility and price are significantly positive indicating that order size—as a fraction of average daily volume—does not decrease with increasing volatility and price as fast as predicted by the invariance hypothesis. The statistically significant positive coefficient on volume may be partially offset by statistically significant negative coefficient on turnover rate.

Discussion. The documented log-normality of bet size is strikingly different from typical assumptions of microstructure models, where innovations in order flow from noise traders are distributed as a normal, not a log-normal or power law. Although normal random variables are a convenient modeling device—they allow conditional expectations to be linear functions of underlying jointly normally distributed variables—their implications are qualitatively very different.

The log-variance of 2.53 of bet size distribution implies that a large fraction of trading volume and even larger fraction of returns variance come from large bets.

Let $\eta(z)$ and $N(z)$ denote the PDF and CDF of a standardized normal distribution, respectively. Define $F(\bar{z}, p)$ by $F(\bar{z}, p) = \int_{z=\bar{z}}^{\infty} \exp(p \cdot \sqrt{2.53} \cdot z) \cdot \eta(z) \cdot dz$. It is easy to show that $F(\bar{z}, p) = \exp(p^2 \cdot 2.53/2) \cdot (1 - N(\bar{z} - p \cdot \sqrt{2.53}))$. This implies that the fraction of the p th moment of order size arising from bets greater than \bar{z} standard deviations above the log-mean is given by $F(\bar{z}, p)/F(-\infty, p) = 1 - N(\bar{z} - p \cdot \sqrt{2.53})$.

Plugging $p = 1$, we find that bets larger than \bar{z} standard deviations above the log-mean (median) generate a fraction of total trading volume given by $1 - N(\bar{z} - \sqrt{2.53})$. Bets larger than the 50th percentile generate 94.41% of trading volume ($\bar{z} = 0$). Bets larger than than $\sqrt{2.53}$ standard deviations above the log-mean (median) bet size—i.e., the largest 5.39% of bets—generate 50% of trading volume ($\bar{z} = \sqrt{2.53}$).

Plugging $p = 2$, we find that bets larger than \bar{z} standard deviations above the log-mean bet size contribute a fraction of total returns variance given by $1 - N(\bar{z} - 2 \cdot \sqrt{2.53})$ under the assumption that the contribution of bets to price variance is proportional to their squared size. Bets greater than the 50th percentile generate 99.93% of returns variance ($\bar{z} = 0$). Bets larger than than $2 \cdot \sqrt{2.53} = 3.18$ standard deviations above the log-mean—i.e., the largest 0.07% of bets—generate 50% of returns variance ($\bar{z} = 2 \cdot \sqrt{2.53}$). For example, if the benchmark stock has about 85 bets per day for each of 252 trading days in a calendar year, the estimates then imply that the 1,155 largest bets out of 21,420 bets during the year generate half of trading volume and approximately the 15 largest bets generate half of returns variance during that year.

Actual trading volume and returns volatility fluctuate over time. Rare large bets may not only account for a significant percentage of returns variance but may also account for some of the stochastic time series variation in volatility. We conjecture that the pattern of short term volatility associated with execution of rare large bets

may depend on the speed with which such bets are executed. Large market disturbances such as the stock market crash of 1929 and 1987, the liquidation of Jerome Kerviel’s trades for Société Générale, and the flash crash of May 6, 2010 could have been induced by execution of gigantic bets.

Another implication of log-normality may be greater kurtosis in the empirical distribution of price changes than a normal distribution would suggest. Given the estimated log-variance of 2.53, the excess kurtosis of one bet has the enormous value of $\exp(10)$ or about 22,000. Thus, excess kurtosis in daily price changes may be influenced more by the kurtosis of individual bets than by the random number of bets arriving each day.

Our thinking about trading games is different from that of the “time change” literature, which goes back to Mandelbrot and Taylor (1967) and Clark (1973). Mandelbrot and Taylor (1967) begin with the intuition that the distribution of price changes is a stable distribution, i.e., a distribution such that a linear combination of two independent random variables has the same shape, up to location and scale parameters. Since it has fatter tails than a normal distribution, it is confined to be a stable Pareto distribution. Following that line of research, the econophysics literature such as Gopikrishnan et al. (1998), Plerou et al. (2000), and Gabaix et al. (2006) estimates different power-laws for the probability distributions of different variables and search for price-formation models consistent with those distributions. Whether order size follows a power law or a log-normal distribution is an interesting question for future research.

Clark (1973) suggests an alternative hypothesis that the distribution of daily price changes is subordinated to a normal distribution with a time clock linked to a log-normally distributed trading volume. The log-normal distribution is neither stable nor infinitely divisible; the sum of random variables with independent log-normal distributions is not log-normal. Thus, if daily price changes can be described by Clark’s hypothesis, neither half-day price changes nor weekly price changes will be described by the same hypothesis.

In some sense, our approach seems to be closer to Mandelbrot and Taylor (1967) who imagine orders of different sizes arriving in the market, with business time linked to their arrival rates rather than trading volume.

The log-normality of bet size may be related to log-normality of assets under management for financial firms. Schwarzkopf and Farmer (2010) study the size of U.S. mutual funds and find that its distribution closely conforms a log-normal with log-variance of about 2.50, similar to our estimates of log-variance for portfolio transition orders. The annual estimates are very stable during twelve years from 1994 to 2005, ranging from 2.43 to 2.59. For years 1991, 1992, and 1993, the log-variance estimates of 1.51, 1.98, and 2.09 are slightly lower, probably because many observations are missing from the CRSP U.S. mutual funds dataset for those years.

Empirical regularities similar to those implied by invariance can be found masked in the previous literature. Bouchaud, Farmer and Lillo (2009) report, for example,

that the number of TAQ prints per day is proportional to market capitalization raised to powers between 0.44 to 0.86. Under the assumption that volatility and turnover rates are stable across stocks as shown in table 1, the midpoint 0.65 of that interval is close to $2/3$ implied by invariance for the number of bets per day, i.e., the inverse of average order size to volume ratio.

6 Empirical Tests Based on Transactions Costs

To examine statistically whether transactions costs conform to the predictions of market microstructure invariance in equation (9), we use the concept of implementation shortfall developed by Perold (1988). Specifically, we estimate costs by comparing the average execution prices of portfolio transition orders with closing prices the evening before any portfolio transition orders begin to be executed. Our tests measure implicit transactions costs resulting from bid-ask spreads and market impact; they exclude explicit transactions costs such as commissions and fees.

Portfolio Transitions and Implementation Shortfall. In portfolio transitions, quantities to be traded are known precisely before trading begins, these quantities are recorded accurately, and all intended quantities are executed. In other trading situations, quantities intended to be traded may not be recorded accurately, and orders may be canceled or quantities may be revised in response to price movements after trading begins. When orders are canceled after prices move in an unfavorable direction or when order size is increased after prices move in a favorable direction, implementation shortfall may dramatically underestimate actual transactions costs. Portfolio transitions data are not subject to these concerns.

Portfolio transition trades are unlikely to be based on short-lived private information about specific stocks because decisions to undertake portfolio transitions and their timing likely result from regularly scheduled meetings of investment committees and boards of plan sponsors, not from fast-breaking private information in the hands of fund managers. Transactions cost estimates are therefore unlikely to be biased upward as a result of short-lived private information being incorporate into prices while orders are being executed.

These properties of portfolio transitions are not often shared by other data. Consider a dataset built up from trades by a mutual fund, a hedge fund, or a proprietary trading desk at an investment bank. In such samples, the intentions of traders may not be recorded in the dataset. For example, a dataset might time stamp a record of a trader placing an order to buy 100,000 shares of stock but not time stamp a record of the trader's actual intention to buy another 200,000 shares after the first 100,000 shares are bought. Furthermore, trading intentions may not coincide with realized trades because the trader changes his mind as market conditions change. Indeed, traders often condition their trading strategies on prices by using limit orders

or canceling orders, thus hard-wiring into their strategies a selection bias problem for using such data to estimate transactions costs. The dependence of actually traded quantities on prices usually makes it impossible to use implementation shortfall in a meaningful way to estimate market depth and bid-ask spreads from data on trades only. Portfolio transitions data are particularly well suited for using implementation shortfall to measure transactions costs because portfolio transitions data avoid these sources of statistical bias.

Non-linear Regression Framework. The predictions invariance makes about transactions costs can be expressed in terms of a non-linear regression. To justify nonlinear regression estimation, we can think of implementation shortfall as representing the sum of two components: (1) the transactions costs incurred as a result of order execution and (2) the effect of other random price changes between the time the benchmark price is set and the time the trades are executed. If we make the identifying assumption that the implementation shortfall from the portfolio transition dataset is an unbiased estimate of the transactions cost, we can think of modeling the other random price changes as an error in a regression of implementation shortfall on transactions costs.

For example, suppose that while one portfolio transition order is being executed, there are 99 other bets being executed at the same time. The temporary and permanent price impact of executing the portfolio transition order shows up as a transactions cost, while the temporary and permanent price impact of the other 99 unobserved bets being executed shows up as other random price changes. Since the portfolio transition order is one of 100 bets being simultaneously executed, the R^2 of the regression is likely to be about 0.01.

For each transition order i , let $\mathbb{I}_{BS,i}$ denote a buy-sell indicator variable which is +1 for buy orders and -1 for sell orders. For transition order i , let C_i denote the transactions cost as a fraction of the value transacted. Let S_i denote the implementation shortfall, defined by $S_i = \mathbb{I}_{BS,i} \cdot (P_{ex,i} - P_{0,i})/P_{0,i}$, where $P_{ex,i}$ is the average execution price of order i and $P_{0,i}$ is its benchmark price. Implementation shortfall is positive when orders are unusually costly and negative when orders are unusually cheap. Since we want the error in our regression be positive when the stock price is moving up and negative when the stock price is moving down, we multiply both the implementation shortfall S_i and the transactions cost $C(X_i)$ by the buy-sell indicator. The regression specification can be then written as $\mathbb{I}_{BS,i} \cdot S_i = \mathbb{I}_{BS,i} \cdot C_i + \tilde{\epsilon}$. Note that $\mathbb{I}_{BS,i} \cdot S_i = (P_{ex,i} - P_{0,i})/P_{0,i}$ since $\mathbb{I}_{BS,i}^2 = 1$.

Invariance imposes the restriction that the unobserved transactions cost C_i has the form given in equation (9), which can be written

$$C_i = \iota^2 \cdot \bar{C}_B \cdot \bar{\sigma}_i \cdot \bar{W}_i^{-1/3} \cdot f\left(\frac{\bar{W}_i^{2/3}}{\iota} \cdot \frac{X_i}{\bar{V}_i}\right). \quad (46)$$

Using equation (1) and equation (2), we replace the “bar” variables $\bar{\sigma}$, \bar{V} , and \bar{W}

with observable variables σ , V , and W and with potentially unobservable constants. This gives us the nonlinear regression

$$\mathbb{I}_{BS,i} \cdot S_i = \mathbb{I}_{BS,i} \cdot \left[\frac{\psi}{\psi^*} \right]^{2/3} \left[\frac{\zeta}{\zeta^*} \right]^{1/3} \left[\frac{\sigma_i}{\sigma^*} \right] \left[\frac{W_i}{W^*} \right]^{-1/3} \cdot f(I_i \cdot \delta^{-1}) / L^* + \tilde{\epsilon}_i, \quad (47)$$

where $I_i := \phi^{-1} \cdot (W_i/W^*)^{2/3} \cdot X_i/V_i$ with invariant constant $\phi := \delta^{-1} \iota \psi^{-2/3} (\zeta/2)^{-1/3} (W^*)^{-2/3}$ obtained from equation (6) and illiquidity measure for the benchmark stock $1/L^* := \iota^2 \cdot \bar{C}_B \cdot \bar{\sigma}^* \cdot [\bar{W}^*]^{-1/3}$ obtained from equation (8). Note that $f(I_i \cdot \delta^{-1})/L^*$ denotes the invariant cost function for the benchmark stock, expressed as a fraction of notional value, similar to equation (9).

Since W_i , X_i , and V_i are observable, the quantity $\phi \cdot I_i$ is observable. The quantity I_i itself in equation (47), however, is not observable because the constant ϕ is defined in terms of potentially unobservable constants ι , δ , ψ , and ζ . To estimate the nonlinear regression equation (47), we substitute for $f(\cdot)$ a different function $f^*(\cdot)$ defined by $f^*(x) = (\psi/\psi^*)^{2/3} \cdot (\zeta/\zeta^*)^{1/3} \cdot f(\phi^{-1} \delta^{-1} x)$. Using $x = \phi \cdot I_i$, the right side of equation (47) becomes a simpler expression in terms of observable data, with various potentially unobserved constants incorporated into the definition of f^* , whose functional form is to be estimated from the data. Under the identifying assumptions $\psi = \psi^*$ and $\zeta = \zeta^*$, we have $f^*(\phi I_i) = f(I_i \cdot \delta^{-1})$. The unobserved constants affect the economic interpretation of the scaling of the estimated functional form for f^* , but they do not otherwise affect the estimation itself.

In order to estimate equation (47), we make two further adjustments, one based on statistics and one based on economics.

First, the variance of errors in the regression is likely to be proportional in size to the variance of returns and the execution horizon. To correct for heteroscedasticity resulting from differences in return volatility, we divide both the right and left sides by return volatility σ_i/σ^* , where $\sigma^* = 0.02$. The root mean squared error of the regression is approximately equal to 0.02, consistent with the interpretation that portfolio transition orders are executed in about one day. Invariance suggests that orders might be executed over horizons inversely proportional to the speed of business time γ , implying very slow executions for large orders in stock with low trading activity. Portfolio transitions are, however, usually implemented within a clearly defined tight time frame, which has the effect of speeding up the “natural” execution horizon for inactive stocks.

Second, to control for the economically and statistically significant influence that general market movements have on implementation shortfall, we add the CRSP value-weighted market return $R_{mkt,i}$ on the first day of the transition to the right side of the regression equation. To the extent that portfolio transition orders are sufficiently large to move the entire U.S. stock market, this adjustment will result in understated transactions costs by measuring only the idiosyncratic component of transactions costs. It is an interesting subject for future research to investigate how large trades in multiple stocks affect general market movements.

Upon making these two changes and using the definition of $f^*(.)$, regression equation (47) becomes

$$\mathbb{I}_{BS,i} \cdot S_i \cdot \frac{(0.02)}{\sigma_i} = \beta_{mkt} \cdot R_{mkt,i} \cdot \frac{(0.02)}{\sigma_i} + \mathbb{I}_{BS,i} \cdot \left[\frac{W_i}{W^*} \right]^\alpha \cdot f^*(\phi I_i)/L^* + \epsilon_i. \quad (48)$$

where invariance implies $\alpha = -1/3$. The function $f^*(\phi I)/L^*$ measures the transactions cost for the benchmark stock in terms of the observable value $\phi I = \tilde{X}/V$. We do not undertake separate estimates of transactions cost parameters for internal crosses, external crosses, and open market transactions. Such estimates would be difficult to interpret due to selection bias resulting from transition managers optimally choosing trading venues to minimize costs.

Although invariance itself does not specify a function form for $f^*(.)/L^*$, the regression places strong cross-sectional restrictions on the shape of the transactions cost function. In addition to the restriction $\alpha = -1/3$, it requires that the same function $f^*(\phi I)/L^*$ with $\phi I = (W_i/W^*)^{2/3} \cdot X_i/V_i$ for order i be used for all stocks.

To adjust standard errors for positive contemporaneous correlation in returns, the observations are pooled by week over the 2001-2005 period into 4,389 clusters across 17 industry categories using the pooling option on Stata.

Dummy Variable Regression. Before reporting regression results for specific functional forms for $f^*(.)$, we first report results of dummy variable regressions of transactions costs.

We sort all 439,765 orders into 100 order size bins of equal size based on the value of “invariant” order size $\phi \cdot I_i = [W_i/W^*]^{2/3} \cdot [X_i/V_i]$. As before, we also place each order into one of ten volume groups based on average dollar trading volume in the underlying stock $P_i V_i$, with thresholds corresponding to the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of NYSE dollar volume. As shown in section 5, the distribution of ϕI_i is approximately invariant across volume groups; across all volume groups $k = 1, \dots, 10$, bin h therefore has a similar number of observations and values of ϕI_i of similar magnitude.

In the regression equation (48), we replace the function $f^*(\phi I_i)/L^*$ with 1,000 dummy variables $D_i^*(k, h)$, $k = 1, \dots, 10$ and $h = 1, \dots, 100$, where $D_i^*(k, h) = 1$ if bet i belongs to volume group k based on dollar volume $P_i V_i$ and to order size bin h based on ϕI_i ; otherwise $D_i^*(k, h) = 0$. We then estimate 1,000 coefficients $f^*(k, h)/L^*$, $k = 1, \dots, 10$, $h = 1, \dots, 100$ for the dummy variables using a separate OLS regression for each of the volume groups, $k = 1, \dots, 10$,

$$\mathbb{I}_{BS,i} \cdot S_i \cdot \frac{(0.02)}{\sigma_i} = \beta_{mkt} \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} + \mathbb{I}_{BS,i} \cdot \left[\frac{W_i}{W^*} \right]^{-1/3} \cdot \sum_{h=1}^{100} D_i^*(k, h) \cdot f^*(k, h)/L^* + \tilde{\epsilon}_i. \quad (49)$$

For volume group k , the dummy variable coefficients $f^*(k, h)/L^*$, $h = 1, \dots, 100$, track the shape of function $f^*(.)/L^*$, without imposing any particular restrictions

on its functional form. Invariance predicts that the ten values of the coefficients $f^*(k, h)/L^*$, $k = 1, \dots, 10$ should be the same for each order size bin h , $h = 1, \dots, 100$.

Figure 4 shows ten plots, one for each of the ten volume groups, with the 100 estimated coefficients for the dummy variables plotted as solid dots in each plot. On each plot, we also superimpose the 95th percent confidence intervals for 100 dummy variable coefficients estimated based on the pooled sample (dotted lines). The superimposed confidence bands help to assess the degree of similarity between cost functions estimated separately based on observations in each volume bin.

On each of the ten plots, the horizontal and vertical axes are scaled in the same way to facilitate comparison. On the horizontal axis, we plot the value for order-size bin h equal to the log of the average ϕI_i for observations in that size bin and corresponding volume group k .

On the right vertical axis, we plot the values of the dummy variable coefficients $f^*(k, h)/L^*$ quantifying for the benchmark stock the cost function as a fraction of notional value, scaled in basis points. To make deviations of cost patterns from invariants visually obvious, we have effectively scaled cost functions as suggested by invariance using regression (49): We multiply orders sizes X_i/V_i by $(W_i/W^*)^{2/3}$ and divide implementation shortfalls S_i by $L^*/L = (\sigma_i/\sigma^*) \cdot (W_i/W^*)^{-1/3}$. Here $1/L^* := \iota^2 \cdot \bar{C}_B \cdot \bar{\sigma}^* \cdot [\bar{W}^*]^{-1/3}$ is the illiquidity measure for the benchmark stock from equation (47).

On the left vertical axis, we plot actual average transactions cost $f^*(k, h)/L^k$ as a fraction of notional value, scaled in basis points. For each volume group k , this scaling reverses invariance-based scaling by multiplying estimated coefficients $f^*(k, h)/L^*$ by L^*/L^k , where $1/L^k$ is the illiquidity measure for orders in volume group k given by $1/L^k := \iota^2 \cdot \bar{C}_B \cdot \bar{\sigma}_{med}^k \cdot (\bar{W}_{med}^k)^{-1/3}$, with $\bar{\sigma}_{med}^k$ denoting median bet volatility and \bar{W}_{med}^k denoting median bet activity for volume group k .

Without appropriate scaling, the data do not reveal its invariant properties. The actual costs on the left vertical axes vary significantly across volume groups. In the low volume group, costs range from -220 basis points to 366 basis points; in the high volume group, costs range from -33 basis point to 55 basis points, 7 times less than in the low volume group.

After applying ‘‘invariance’’ scaling, however, our plots appear to be visually consistent with the invariance hypothesis. For all ten subplots in figure 4, the estimated dummy variable coefficients on the right vertical axes are very similar across volume groups. They also line up along the superimposed confidence band.

Moving from low-volume groups to high-volume groups, these estimates also become visually more noisy. For low volume group one, dummy estimates lie within the confidence band, very tightly pinning down the estimated shape for the function $f^*(\cdot)/L^*$. For volume group ten, many dummy estimates lie outside of the confidence band, with 11 observations above the band and about 40 observations being below. These patterns suggest that the statistical power of our tests concerning transactions costs comes mostly from low-volume groups.

Since transition orders are executed over a fixed and limited number of calendar days, the execution in business time is faster for low-volume stocks and slower for high-volume stocks. When a transition order in a low-volume stock is being executed, there are therefore probably fewer other bets being executed at the same time; this makes the R^2 of the regression higher. In the same amount of calendar time, many more bets are being executed for the high-volume stocks, making the R^2 much lower. The more patient business-time pace of execution for high-volume stocks may explain why their execution costs appear to be slightly less expensive than low-volume stocks.

Transactions Cost Estimates in Non-Linear Regression. We next report estimates of specific parametric functional forms for cost function $f^*(\cdot)/L^*$, assumed to be the sum of a constant bid-ask spread term and a market impact term which is a power of $\phi \cdot I$. For this particular specification, the non-linear regression (48) can be written as

$$\mathbb{I}_{BS,i} \cdot S_i \cdot \frac{(0.02)}{\sigma_i} = \beta_{mkt} \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} + \mathbb{I}_{BS,i} \cdot \bar{\kappa}_0^* \cdot \left[\frac{W_i}{W^*} \right]^{\alpha_1} + \mathbb{I}_{BS,i} \cdot \bar{\kappa}_I^* \cdot \left[\frac{W_i}{W^*} \right]^{\alpha_2} \cdot \left[\frac{\phi I_i}{0.01} \right]^z + \tilde{\epsilon}. \quad (50)$$

where $\phi I_i/0.01 = (W_i/W^*)^{2/3} \cdot \frac{X_i}{(0.01)V_i}$. The explanatory variables are scaled so that, for the benchmark stock, execution of one percent of daily volume has price impact cost of $\bar{\kappa}_I^*$ and fixed bid-ask spread of $\bar{\kappa}_0^*$, both measured as a fraction of the value traded, with units of 10^{-4} equivalent to basis points in this paper. Equation (50) is an empirical version of equation (41).

First, we report estimates of the six parameters $(\beta_{mkt}, z, \alpha_1, \bar{\kappa}_0^*, \alpha_2, \bar{\kappa}_I^*)$ in equation (50) using non-linear regression. Second, we calibrate the three-parameter linear impact model of equation (11) with parameters $(\beta_{mkt}, \bar{\kappa}_0^*, \bar{\kappa}_I^*)$ by imposing the additional invariance restrictions $\alpha_1 = \alpha_2 = -1/3$ and the linear cost restriction $z = 1$. Third, we also calibrate the three-parameter square root model of equation (12) with parameters $(\beta_{mkt}, \bar{\kappa}_0^*, \bar{\kappa}_I^*)$ by imposing the alternate restriction $z = 1/2$. Finally, we examine a twelve-parameter generalization of equation (50) which replaces powers α_1 and α_2 of trading activity W with powers of volatility σ , price P , volume V , and turnover η . The results support invariance, with the square root version of invariance explaining transactions costs somewhat better than the linear version.

The parameter estimates for the six parameters $\beta_{mkt}, \bar{\kappa}_0^*, z, \alpha_1, \bar{\kappa}_I^*, \alpha_2$ in the non-linear regression (50) are reported in table 4.

For the coefficient β_{mkt} which multiplies the market return R_{mkt} , the estimate is $\hat{\beta}_{mkt} = 0.65$ with standard error 0.013. The fact that $\hat{\beta}_{mkt} < 1$ suggests that many transition orders are executed early on the first day.

The point estimate of the estimated bid-ask spread exponent is $\hat{\alpha}_1 = -0.49$, with standard error 0.050, three standard errors lower than the predicted value $\alpha_1 = -1/3$. This result may not be economically meaningful because the estimated coefficient on the bid-ask spread itself, $2 \cdot \bar{\kappa}_0^*$, is not significantly different from zero.

The point estimate for α_2 is $\hat{\alpha}_2 = -0.32$ with standard error 0.015. Since invariance implies $\alpha_2 = -1/3$, this result strongly supports invariance. When the four parameters are estimated separately for NYSE Buys, NYSE Sells, NASDAQ Buys, and NASDAQ Sells, the estimated coefficient for α_2 are -0.40 , -0.33 , -0.41 , and -0.29 , respectively.

The estimate for the market impact curvature parameter z is $\hat{z} = 0.57$ with standard error 0.039. This suggests that a square-root specification ($z = 1/2$) may describe observed transactions costs better than a linear specification ($z = 1$). Note that invariance does not place any restrictions on parameter z itself.

The point estimate of $\bar{\kappa}_I^*$ is $\hat{\kappa}_I^* = 10.69 \cdot 10^{-4}$ with standard error $1.376 \cdot 10^{-4}$. The estimates of $\bar{\kappa}_I^*$ are higher for buy orders than for sell orders ($12.08 \cdot 10^{-4}$ versus $9.56 \cdot 10^{-4}$ for NYSE; $12.33 \cdot 10^{-4}$ versus $9.34 \cdot 10^{-4}$ for NASDAQ).

Results for the five-parameter linear specification, regression equation (50) with parameter z restricted to be equal to one, are reported in table 9 in the Appendix. The estimate of the bid-ask spread cost $\bar{\kappa}_0^*$ is $6.28 \cdot 10^{-4}$ with standard error $0.890 \cdot 10^{-4}$ and the estimate of the exponent α_1 is $\hat{\alpha}_1 = -0.39$ with standard error 0.020. The estimate of the market impact cost $\bar{\kappa}_I^*$ is $2.73 \cdot 10^{-4}$ with standard error 0.252, and the estimate of the exponent α_2 is $\hat{\alpha}_2 = -0.31$ with standard error of 0.028.

Model Calibration. Table 5 presents estimates for the three parameters β_{mkt} , $\bar{\kappa}_0^*$, and $\bar{\kappa}_I^*$ in equation (50), imposing the invariance restrictions $\alpha_1 = \alpha_2 = -1/3$ and also imposing either a linear transactions cost model $z = 1$ or a square root model $z = 1/2$.

In the linear specification with $z = 1$, the point estimate for market impact cost $\hat{\kappa}_I^*$ is equal to $2.5003 \cdot 10^{-4}$, and the point estimate for bid-ask spread cost $\hat{\kappa}_0^*$ is equal $8.2134 \cdot 10^{-4}$. For the benchmark stock, these estimates imply that the total cost of a hypothetical trade of one percent of daily volume incurs the costs of about 10.71 basis points.

In the square-root specification with $z = 1/2$, the point estimate for market impact cost $\hat{\kappa}_I^*$ is equal to $12.0787 \cdot 10^{-4}$, and the point estimate for half bid-ask spread $\hat{\kappa}_0^*$ is equal $2.0763 \cdot 10^{-4}$.

The benchmark stock belongs to volume group seven, and the corresponding average quoted spread in table 1 for that group is 12.04 basis points. The implied spread estimate of about 16.42 basis points for the linear model is close to the quoted spread, whereas the spread estimate of 4.16 basis points for the square root model could have been biased downwards due to collinearity between a constant term and a square-root term in the regression in a region next to zero.

Economic Interpretation. A comparison of the R-squares in table 4 and table 5 provides strong support for the invariance. When the coefficient on W/W_* is fixed at the invariance-implied value of $-1/3$ and only two transactions cost parameters $\bar{\kappa}_I^*$

and $\bar{\kappa}_0^*$ are estimated, as in table 5, the R-square is $R^2 = 0.0991$ for a linear specification and $R^2 = 0.1007$ for a square-root specification. The square root specification performs slightly better than the linear specification. Adding the three additional parameters α_1, α_2 and z only mildly increases the R-square to $R^2 = 0.1010$, as in table 5.

We also consider a more general specification with eleven coefficients estimated. The exponents on the three components of trading activity W_i (volatility σ_i , price $P_{0,i}$, volume V_i) as well as the coefficient on the monthly turnover ν_i are allowed to vary freely. The estimated regression equation is

$$\mathbb{I}_{BS,i} \cdot S_i \cdot \frac{(0.02)}{\sigma_i} = \beta_{mkt} \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} \mathbb{I}_{BS,i} \cdot \bar{\kappa}_0^* \cdot \left[\frac{W_i}{W^*} \right]^{-1/3} \cdot \frac{\sigma_i^{\beta_1} \cdot P_{0,i}^{\beta_2} \cdot V_i^{\beta_3} \cdot \nu_i^{\beta_4}}{(0.02)(40)(10^6)(1/12)} + \\ + \mathbb{I}_{BS,i} \cdot \bar{\kappa}_I^* \cdot \left[\frac{\phi I_i}{0.01} \right]^z \cdot \left[\frac{W_i}{W^*} \right]^{-1/3} \cdot \frac{\sigma_i^{\beta_5} \cdot P_{0,i}^{\beta_6} \cdot V_i^{\beta_7} \cdot \nu_i^{\beta_8}}{(0.02)(40)(10^6)(1/12)} + \tilde{\epsilon}.$$

Because the exponents on the W -terms are set to be $-1/3$, the invariance hypothesis predicts $\beta_1 = \beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = \beta_7 = \beta_8 = 0$.

Table 5 shows that despite increasing the number of estimated parameters from four to eleven, the R-square in the aggregate regression increases from $R^2 = 0.1010$ to only $R^2 = 0.1016$. The estimates of $\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7$, and β_8 are shown in table 11 in Appendix. The estimates of $\beta_1, \beta_2, \beta_3, \beta_4$ are often statistically significant, but these explanatory variables are multiplied by statistically insignificant coefficient $\bar{\kappa}_0^*$. Almost all estimates of $\beta_5, \beta_6, \beta_7$, and β_8 are statistically insignificant, both for the pooled sample as well as four sub-samples.

In all three specifications, separate regressions for NYSE Buys, NYSE Sells, NASDAQ Buys, and NASDAQ Sells suggest that price impact costs are higher for buy orders than for sell orders. This is consistent with the idea in Obizhaeva (2009) that the market believes that buy orders—especially buy orders in portfolio transitions—contain more information than sell orders.

OLS Estimates for Quoted Spread. Since invariance implies that bid-ask spread costs are proportional to $\bar{\sigma} \cdot \bar{W}^{-1/3}$, intuition suggests that quoted spreads may also have this invariant property. As a supplement to our empirical results on transactions costs, we test this prediction using data on quoted spreads, supplied in the portfolio transition data as pre-trade information for each transition order.

Let s_i denote the dollar quoted spread for order i . Applying equation (11) or (12) to quoted spreads, we obtain

$$s_i/P_i \propto \bar{\sigma}_i \cdot \bar{W}_i^{-1/3}. \quad (51)$$

Using equation (1) and equation (2), we can write the log-linear OLS regression

$$\ln \left[\frac{s_i}{P_i \cdot \sigma_i} \right] = \ln \bar{s} + \alpha_3 \cdot \ln \left[\frac{W_i}{W^*} \right] + \tilde{\epsilon}_i, \quad (52)$$

where invariance implies $\alpha_3 = -1/3$. The constant term $\ln \bar{s} := \ln[s^*/(40 \cdot 0.02)] + 2/3 \ln(\psi/\psi^*) - 1/3 \ln(\zeta^*/\zeta)$ quantifies the dollar spread s^* for the benchmark stock as a fraction of dollar volatility $P^* \cdot \sigma^*$, under the identifying assumptions $\zeta = \zeta^*$ and $\psi = \psi^*$.

Table 6 presents the regression results. The point estimate $\hat{\alpha}_3 = -0.35$ has standard error 0.003. For sub-samples of NYSE Buys, NYSE Sells, NASDAQ Buys, and NASDAQ Sells, the estimates are -0.31 , -0.32 , -0.40 , and -0.39 , respectively. Although the hypothesis $\alpha_3 = -1/3$ is usually rejected, the estimates are economically very close to $-1/3$ predicted by invariance. The point estimate of $\ln \bar{s}$ is equal to -3.07 , implying a quoted spread of $\exp(-3.07) \cdot 0.02 \approx 9 \cdot 10^{-4}$ for the benchmark stock. This number is similar to median spread of 8.12 basis points for volume group seven in table 1.

It can be shown that an implicit spread proportional to $\bar{\sigma} \cdot \bar{W}^{-1/3}$, as implied by invariance, provides a better proxy for the actually incurred spread costs than the quoted spread itself. If regression equation (50) is estimated with linear impact $z = 1$, using only the 436,649 observations for which quoted bid-ask spread data is supplied, we obtain $R^2 = 0.0992$. Now replace the invariance implied spread cost proportional to $\bar{\sigma} \cdot \bar{W}^{-1/3}$ with the quoted half spread $1/2 \cdot s_i/P_i$ in equation (11). The estimated equation is

$$\mathbb{I}_{BS,i} \cdot S_i \cdot \frac{(0.02)}{\sigma_i} = \beta_{mkt} \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} + \mathbb{I}_{BS,i} \cdot h \cdot \frac{1}{2} \cdot \frac{s_i}{P_{0,i}} \cdot \frac{(0.02)}{\sigma_i} + \mathbb{I}_{BS,i} \cdot \bar{\kappa}_I \cdot \left[\frac{\phi I_i}{0.01} \right] \cdot \left[\frac{W_i}{W^*} \right]^{\alpha_2} + \tilde{\epsilon}.$$

We find that the R^2 drops from $R^2 = 0.0992$ to $R^2 = 0.0976$ (see table 10 in the Appendix). The point estimate of the coefficient on the quoted half-spread coefficient is $\hat{h} = 0.71$. The estimates are equal to 0.61, 0.74, 0.61, and 0.75, when estimated for NYSE Buys, NYSE Sells, NASDAQ Buys, and NASDAQ Sells. One interpretation of the estimate of 0.71 is that transition managers incur as a transactions cost only 71% of the quoted half-spread. The values are consistent with the intuition in Goettler, Parlour and Rajan (2005) that endogenously optimizing traders capture a fraction of the bid-ask spread by mixing between market orders and limit orders. It is also possible that noise in the quoted spread biases the coefficient towards zero and reduces the explanatory power of the regression. Note that this errors in variables problem does not bias the coefficient estimate in equation (52), where the quoted spread is on the left side.

Consistent with our findings and predictions of invariance, Bouchaud, Farmer and Lillo (2009) and Dufour and Engle (2000) report that the quoted bid-ask spread is proportional to the standard deviation of percentage returns between trades.

Discussion. Figure 5 plots the 100 estimated coefficients for the dummy variables and their 95th confidence intervals from regression equation (49) estimated for pooled sample. The linear and square root cost functions with parameters calibrated in table 5 are superimposed. A linear specification is $2.50 \cdot 10^{-4} \cdot \phi I / 0.01 + 8.21 \cdot 10^{-4}$ (black

solid line), and a square-root specification is $12.07 \cdot 10^{-4} \cdot \sqrt{\phi I / 0.01} + 2.08 \cdot 10^{-4}$ (grey solid line). Both specifications result in estimates economically close to each other.

Consistent with the higher reported R^2 for the square root model than the linear model in table 5, the square-root specification fits the data slightly better than the linear specification, particularly for large orders in the size bins from 90th to 99th percentile. Similarly, most studies find that total price impact is best described by a concave function. For example, Almgren et al. (2005) find the estimate of temporary price impact curvature $\hat{z} = 0.60$ for their sample of almost 30,000 U.S. stock orders executed by Citigroup between 2001 and 2003, comparable to our estimate of $\hat{z} = 0.56$ when we fix $\alpha_1 = \alpha_2 = -1/3$ in regression equation (50). To differentiate temporary impact from permanent impact of earlier executed trades, Almgren et al. (2005) assume a particular execution algorithm with a constant rate of trading. We do not quantify these costs components separately but rather focuses on the total costs.

Intuition might suggest that for gigantic orders, the square root model would predict dramatically lower transactions costs than the linear model, making it easy to distinguish the predictions of one model from the other. When the estimated linear and square root cost functions are superimposed on ten plots in figure 4, we find that both specifications provide economically similar estimates of transactions costs, even up to the 99th and 100th size bins. Consistent with our results about bet size, there are not enough observations of gigantic transition orders in high-volume stocks to create statistically compelling differences between the square root and linear models for the most active stocks. For the very largest orders in the 100th size bin for the highest volume group in figure 4, the estimated dummy variable fits the higher cost estimates of the linear model better than the square root model. In comparison with the square root model, this suggests that linear specification provides better predictions of market impact for very large orders in very active markets.

7 Implications.

The invariance relationships (5), (6), (9), (13), and (14) are like a structural model which describes how the microstructure of financial markets work. The model is fully specified by constants describing the moments of \tilde{I} and the shape of the unmodeled function $C_B(\cdot)$, which determines \bar{C}_B . These constants can be inferred from the estimates in section 5 and section 6, but their interpretation depends on the assumptions about volume multiplier ζ , volatility multiplier ψ , and deflator δ .

Our empirical tests provide not only evidence in favor of the invariance hypotheses but also inputs for the calibration process. Our empirical results can be summarized as follows. The distribution of portfolio transition orders $|\tilde{X}|$ —expressed as a fraction of volume—is approximately a log-normal. It is therefore fully described by two parameters, the log-mean for the benchmark stock estimated to be -5.71 and the log-variance estimated to be 2.53 (see table 3). The following formula shows how these estimates can be extrapolated to stocks with other levels of trading activity

$W = \sigma PV$ and volume V :

$$\ln \left[\frac{|\tilde{X}|}{V} \right] \approx -5.71 - \frac{2}{3} \cdot \ln \left[\frac{W}{(0.02)(40)(10^6)} \right] + \sqrt{2.53} \cdot \tilde{Z}, \quad \tilde{Z} \sim N(0, 1) \quad (53)$$

Our empirical results also suggest that transactions cost functions can be described by either a linear model or a square root model. Since both models also have a constant bid-ask spread term, each model is described by two parameters. For an order of 1% of average daily volume in the benchmark stock, the estimates imply market impact costs of $\bar{\kappa}_I = 2.50 \cdot 10^{-4}$ and spread costs of $\bar{\kappa}_0 = 8.21 \cdot 10^{-4}$ for the linear model as well as market impact costs of $\bar{\kappa}_I = 12.08 \cdot 10^{-4}$ and spread costs of $\bar{\kappa}_0 = 2.08 \cdot 10^{-4}$ for the square root model (see table 5). The following formulas show how these estimates can be extrapolated to execution costs of an order of X shares for stocks with other levels of trading activity W , volume V and volatility σ :

$$C(X) = \frac{\sigma}{0.02} \left(\frac{8.21}{10^4} \cdot \left[\frac{W}{(0.02)(40)(10^6)} \right]^{-1/3} + \frac{2.50}{10^4} \cdot \left[\frac{W}{(0.02)(40)(10^6)} \right]^{1/3} \frac{X}{(0.01)V} \right). \quad (54)$$

$$C(X) = \frac{\sigma}{0.02} \left(\frac{2.08}{10^4} \cdot \left[\frac{W}{(0.02)(40)(10^6)} \right]^{-1/3} + \frac{12.08}{10^4} \cdot \left[\frac{X}{(0.01)V} \right]^{1/2} \right). \quad (55)$$

We calibrate the distribution of the invariant \tilde{I} using calibrated values for the distribution of order sizes in equation (53). The distribution of \tilde{I} is linked to the distribution of portfolio transition orders \tilde{X} by equation $\phi \cdot |\tilde{I}| = \tilde{X}/V \cdot (W/W^*)^{2/3}$ with the constant $\phi := \iota \delta^{-1} \psi^{-2/3} (\zeta/2)^{-1/3} (W^*)^{-2/3}$, as explained in the discussion after equation (47). The log-normality of portfolio transition orders implies the log-normality of $|\tilde{I}|$, and the estimated moments of $\tilde{X}/V \cdot (W/W^*)^{2/3}$ put restrictions on moments of invariant \tilde{I} .

The distribution of invariant \tilde{I} can be therefore approximated by the product of two independent random variables: (1) a buy-sell indicator variable $\tilde{\mathbb{I}}_{BS}$ taking values of +1 or -1 with equal probability, and (2) a log-normal random variable with log-mean μ_I and log-variance σ_I^2 . Matching log-moments of $\phi \cdot |\tilde{I}|$ to calibrated constants of -5.71 and 2.53 from equation (53)— $E\{\ln(\phi \cdot |\tilde{I}|)\} = -5.71$ and $V\{\ln(\phi \cdot |\tilde{I}|)\} = 2.53$ —we find

$$\tilde{I} \sim \tilde{\mathbb{I}}_{BS} \cdot e^{\mu_I + \sigma_I \cdot \tilde{Z}} = \tilde{\mathbb{I}}_{BS} \cdot e^{5.6599 + \ln \psi + \frac{1}{2} \ln(\zeta/2) + \frac{3}{2} \ln \delta + \sqrt{2.53} \cdot \tilde{Z}}. \quad (56)$$

In the baseline case of $\psi = \zeta/2 = \delta = 1$, the invariant $\tilde{I} \sim \tilde{\mathbb{I}}_{BS} \cdot \exp\{5.6599 + \sqrt{2.53} \cdot \tilde{Z}\}$. More generally, interpretation of the estimated constants in equation (53) depends on assumptions about values of the volatility multiplier ψ , the intermediation multiplier ζ , and the deflator δ .

We calibrate the average cost of a bet $\bar{C}_B = E\{C(\tilde{Q}) \cdot P \cdot |\tilde{Q}|\}$ using two approaches. One approach is based on the linear model, and the other approach is based on the square root model. Both approaches assume scaled portfolio transition orders $\delta \cdot X_i$

have the same distribution as a typical bet $|\tilde{Q}|$, implying that $\bar{C}_B = E\{C(\tilde{X} \cdot \delta) \cdot P \cdot \tilde{X} \cdot \delta\}$. First, using the linear model from equations (53) and (54), we obtain

$$\bar{C}_B = 1730 \cdot \delta^2 + 385 \cdot \delta. \quad (57)$$

Second, using the square root model from equations (53) and (55), we obtain

$$\bar{C}_B = 1586 \cdot \delta^{3/2} + 98 \cdot \delta. \quad (58)$$

In the baseline case with $\delta = 1$, the linear model implies cost of executing a bet $\bar{C}_B = \$2,115$, which is higher than the cost $\bar{C}_B = \$1,684$ implied by the square root model.

The calibrated invariants \tilde{I} and \bar{C}_B imply specific quantitative relationships concerning various market microstructure variables. Those predictions depend on our assumptions about volatility multiplier ψ , volume multiplier ζ , and deflator δ .

The structure of the meta-model imposes the additional restriction on the invariants and says that the standard deviation of invariant \tilde{I} is equal to invariant \bar{C}_B , i.e., equivalently $\bar{C}_B = E|\tilde{I}|/m$, as seen from equation (31). Note that this restriction does not appear during the discussion of the invariance hypotheses in section 1. The economic intuition of this restriction is that market makers must be able to break even trading against bets which move prices due to their information content. Using equation (56), we find $m := E\{|\tilde{Q}|\}/\{E\{\tilde{Q}^2\}\}^{1/2} \approx \exp(-\sqrt{2.53}/2) \approx 0.28$ and write the restriction in terms of unidentified parameters of volume multiplier ζ , volatility multiplier ψ , and deflator δ as

$$\bar{C}_B = \$3569 \cdot \delta^{3/2} \cdot (\zeta/2)^{1/2} \cdot \psi, \quad (59)$$

where \bar{C}_B depends on deflator δ as shown either in equation (57) or in equation (58).

The equation above is not satisfied for a baseline set of parameters $\psi = \zeta/2 = \delta = 1$. There exist, however, many parameters satisfying that constraint. For illustration purposes, we select two sets of arbitrary parameters with reasonable values: (1) $\psi = 0.55$, $\zeta = 2.3$, and $\delta = 1$ for the linear model and (2) $\psi = 0.5$, $\zeta = 2$, and $\delta = 0.4$ for the square root model.

For these sets of parameters consistent with the meta-model, the implications of invariance hypotheses can be stated in a consolidated form as in the invariance theorem in section 2. To derive those implications, we substitute \bar{V} , $\bar{\sigma}$ and \bar{W} for V , σ and W in equations (30) and (31), then use equation (1), equation (2) and $m \approx 0.28$ from equation (56),

$$\gamma = \left(\frac{\lambda \cdot (2/\zeta)V}{\psi\sigma \cdot P \cdot m} \right)^2 = \left(\frac{E\{|\tilde{Q}|\}}{(2/\zeta)V} \right)^{-1} = \frac{(\psi\sigma \cdot L)^2}{0.28^2} = \frac{\psi^2\sigma^2}{\theta^2\tau \cdot \Sigma} = \frac{\rho}{\theta^2\tau} = \left(\frac{(2/\zeta)\psi W}{0.28\bar{C}_B} \right)^{2/3}. \quad (60)$$

$$\frac{\tilde{Q}}{\bar{V}} \sim \delta \cdot \left[\frac{W}{(0.02)(40)(10^6)} \right]^{-2/3} \cdot e^{-5.71 + \sqrt{2.53} \cdot \tilde{Z}}, \quad \tilde{Z} \sim N(0, 1). \quad (61)$$

Given estimates of \bar{C}_B and particular choice of the constants ψ , ζ , and δ , these equations summarize the main quantitative implications of the invariance hypotheses, calibrated based on portfolio transition data and consistent with the meta-model.

The values of the number of bets per day γ , their size \tilde{Q} , long-term market impact λ , illiquidity $1/L$, efficiency $(1/\Sigma)^{1/2}$, resilience ρ , and cost functions can be immediately calculated given observable W and its components V and σ . For the benchmark stock with price \$40, volume 1 million shares per day, and volatility 2 percent per day, for example, some of the calibrated variables are presented in the table.

	\tilde{C}_B	γ^*	$E\{ \tilde{Q}^* \}/V^*$	λ^*	$1/L^*$
(1) $\psi = 0.55, \zeta = 2.3, \delta = 1.0$	\$2,115	75	1.16%	$1.23 \cdot 10^{-6}$	$45 \cdot 10^{-4}$
(2) $\psi = 0.50, \zeta = 2.0, \delta = 0.4$	\$440	219	0.46%	$1.67 \cdot 10^{-6}$	$24 \cdot 10^{-4}$

For an order equal to 1% of average daily volume, for example, the calibrated parameter of the long-term impact for linear case implies that the percentage long-term impact $\lambda \cdot 10^6 \cdot 0.01/40$ is equal to $3.08 \cdot 10^{-4}$. Since bid-ask spread costs are estimated to be less than impact costs for the linear model, the linear model implies not only that bid-ask spread costs $8.21 \cdot 10^{-4}$ are transitory, but also a portion of the linear impact $2 \cdot 2.50 \cdot 10^{-4}$ of the “last” trade in a bet will, on average, go away as the price mean-reverts to its long-term level. For square root transactions costs, the bid-ask spread costs and some of the price impact costs will mean revert away as well, but for gigantic bets, the concavity of the square root function implies greater long-term impact than the short-term impact associated with transactions costs. Of course, these implications apply to bets of “average” information content. In a more general model, markets might figure out quickly that some bets are informed and other bets are not; gigantic bets may have information content less than proportional to their sizes.

The decomposition of total impact into permanent and temporary components is closely related to market efficiency and resilience, which depend on the third invariants quantifying the precision per bet $\theta^2\tau$ in equation (30). Calibration of this invariant and also deeper understanding of temporary and permanent impacts require the analysis of price dynamics induced by portfolio transition orders or other transactions. We leave these important questions for the future research.

Our calibration relies on a number of assumptions regarding whether portfolio transitions are similar in terms of size and associated transactions costs to typical bets. In the future, a better calibration of the invariants \tilde{I} , \bar{C}_B , and $\theta^2\tau$ as well as estimation of multipliers ψ , ζ , and δ will be necessary to sharpen predictions based on invariance principles.

8 Conclusion: Future Research

We have shown that the predictions based on market microstructure invariance largely hold in portfolio transitions data for equities. We conjecture that predictions based on invariance can be found to hold in other data as well, such as transactions in the Trades and Quotes (“TAQ”) dataset, changes in holdings recorded in 13-F filings of institutional investment managers, institutional trades reported in the Ancerno dataset, and other datasets. We conjecture that data on news articles can help to show that information flows are confined to the same business time as trading.

We conjecture that predictions of market microstructure invariance may generalize to other markets such as bond markets, currency markets, and futures markets, as well as to other countries. Whether market microstructure invariance applies this generally poses an interesting set of issues for future research.

References

- Almgren, Robert, Chee Thum, Emmanuel Hauptmann, and Hong Li.** 2005. “Direct Estimation of Equity Market Impact.” *Risk*, 18(7): 58–62.
- Amihud, Yakov.** 2002. “Illiquidity and Stock Returns: Cross-Section and Time-Series Effects.” *The Journal of Financial Markets*, 5: 31–56.
- Ané, Thierry, and Helyette Geman.** 2000. “Order Flow, Transaction Clock, and Normality of Asset Returns.” *The Journal of Finance*, 55(5): 2259–2284.
- Back, Kerry, and Shmuel Baruch.** 2004. “Information in Securities Markets: Kyle meets Glosten and Milgrom.” *Econometrica*, 72: 433–465.
- Black, Fischer.** 1986. “Noise.” *The Journal of Finance*, 41: 529–543.
- Bouchaud, Jean-Philippe, Doyne J. Farmer, and Fabrizio Lillo.** 2009. “How Markets Slowly Digest Changes in Supply and Demand.” In *Handbook of Financial Markets: Dynamics and Evolution*. 47–160. North Holland.
- Breen, William, Laurie Simon Hodrick, and Robert A. Korajczyk.** 2002. “Predicting Equity Liquidity.” *Management Science*, 48: 470–83.
- Brennan, Michael, and Avanidhar Subrahmanyam.** 1998. “The Determinants of Average Trade Size.” *The Journal of Business*, 71(1): 1–25.
- Campbell, John Y., and Albert S. Kyle.** 1993. “Smart Money, Noise Trading, and Stock Price Behavior.” *The Review of Economic Studies*, 60: 1–34.
- Clark, Peter.** 1973. “A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices.” *Econometrica*, 41: 135–155.

- Derman, Emanuel.** 2002. “The Perception of Time, Risk and Return during Periods of Speculation.” *Quantitative Finance*, 2: 282–296.
- Dufour, Alfonso, and Robert F. Engle.** 2000. “Time and the Price Impact of a Trade.” *The Journal of Finance*, 55(6): 2467–2598.
- Farmer, Doyme, Austin Gerig, Fabrizio Lillo, and Henri Waelbroeck.** 2012. “How Efficiency Shapes Market Impact.” Working Paper.
- Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and Eugene H. Stanley.** 2006. “Institutional Investors and Stock Market Volatility.” *The Quarterly Journal of Economics*, 121: 461–504.
- Glosten, Lawrence R., and Paul R. Milgrom.** 1985. “Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders.” *The Journal of Financial Economics*, 14: 71–100.
- Goettler, Ronald L., Christine A. Parlour, and Uday Rajan.** 2005. “Equilibrium in a Dynamic Limit Order Market.” *The Journal of Finance*, 60(5): 2149–2192.
- Gopikrishnan, Parameswaran, Martin Meyer, Luis A. Amaral, and Eugene H. Stanley.** 1998. “Inverse Cubic Law for the Distribution of Stock Price Variations.” *The European Physical Journal*, 3: 139–140.
- Grossman, Sanford J., and Joseph E. Stiglitz.** 1980. “On the Impossibility of Informationally Efficient Markets.” *American Economic Review*, 70(3): 393–408.
- Grossman, Sanford J., and Merton H. Miller.** 1988. “Liquidity and Market Structure.” *The Journal of Finance*, 43(3): 617–633.
- Hasbrouck, Joel.** 1999. “Trading Fast and Slow: Security Market Events in Real Time.” Working Paper, New York University.
- Hasbrouck, Joel.** 2007. *Empirical Market Microstructure: The Institutions, Economics, and Econometrics of Securities Trading*. Oxford University Press, USA.
- Kyle, Albert S.** 1985. “Continuous Auctions and Insider Trading.” *Econometrica*, 53(6): 1315–1335.
- Lee, Charles, and Mark Ready.** 1991. “Inferring Trade Direction from Intraday Data.” *The Journal of Finance*, 46: 733–746.
- Mandelbrot, Benoît, and Howard M. Taylor.** 1967. “On the Distribution of Stock Price Differences.” *Operations Research*, 15(6): 1057–1062.
- Modigliani, Franco, and Merton H. Miller.** 1958. “The Cost of Capital, Corporation Finance and the Theory of Investment.” *American Economic Review*, 48: 261–297.

- Obizhaeva, Anna.** 2009. “Portfolio Transitions and Price Dynamics.” Working Paper, University of Maryland.
- Perold, André.** 1988. “The Implementation Shortfall: Paper vs. Reality.” *Journal of Portfolio Management*, 14(3): 4–9.
- Plerou, Vasiliki, Parameswaran Gopikrishnan, Luis A. Amaral, Xavier Gabaix, and Eugene H. Stanley.** 2000. “Economic Fluctuations and Anomalous Diffusion.” *Physical Review E*, 62: 3023–3026.
- Schwarzkopf, Yonathan, and Doyne J. Farmer.** 2010. “Empirical Study of the Tails of Mutual Fund Size.” *Physical Review E*, 81: 066113.
- Stambough, Robert F., and Lubos Pastor.** 2003. “Liquidity Risk and Expected Stock Returns.” *Journal of Political Economy*, 111: 642–685.
- Treynor, Jack.** 1995. “The Only Game in Town.” *The Financial Analysts Journal*, 51(1): 81–83 (reprinted from *The Financial Analysts Journal* 22, 1971, 12–14, 22).
- Wang, Jiang.** 1993. “A Model of Intertemporal Asset Prices Under Asymmetric Information.” *The Review of Economic Studies*, 60(2): 249–282.

Table 1: Descriptive Statistics.

Panel A: Properties of Securities.

	All	1	2	3	4	5	6	7	8	9	10
Med($V \cdot P$) $\times 10^6$	18.72	1.13	5.10	9.92	15.93	23.87	31.41	42.12	60.25	101.60	212.85
Med(σ) $\times 10^2$	1.93	2.16	2.04	1.94	1.98	1.90	1.86	1.80	1.78	1.77	1.76
Med(S_{prd}) $\times 10^4$	12.04	40.96	18.72	13.70	12.02	10.32	9.42	8.12	7.21	5.92	4.83
Med(T_{urn}) $\times 10^2$	12.72	6.58	12.36	14.80	15.95	16.80	16.36	15.37	14.41	14.92	11.12

Panel B: Properties of Portfolio Transitions Orders.

	All	1	2	3	4	5	6	7	8	9	10
Avg(X/V) $\times 10^2$	4.20	16.23	4.54	2.62	1.83	1.37	1.18	1.08	0.88	0.69	0.49
Med(X/V) $\times 10^2$	0.57	3.33	1.36	0.79	0.53	0.40	0.34	0.30	0.25	0.20	0.14
Avg(X/Cap) $\times 10^4$	1.72	3.55	2.68	2.04	1.59	1.26	1.06	0.91	0.72	0.56	0.37
Med(X/Cap) $\times 10^4$	0.35	0.98	0.80	0.58	0.42	0.32	0.27	0.23	0.19	0.15	0.09
Avg $C(X)$ $\times 10^4$	16.79	44.95	21.46	14.53	12.62	11.70	5.58	9.27	3.99	7.37	6.16
Avg Comm $\times 10^4$	7.43	14.90	9.30	7.86	7.00	6.15	5.49	4.93	4.34	3.62	2.68
Avg SEC fee $\times 10^5$	2.90	3.26	3.02	3.00	2.85	2.84	2.76	2.76	2.73	2.68	2.56
# Obs	439,765	71,000	68,689	41,238	49,000	28,073	29,330	29,778	34,409	40,640	47,608
# Stks	2,583	1,108	486	224	182	106	126	90	102	81	78

Table reports the characteristics of securities and transition orders. Panel A shows the median average daily dollar volume (in millions of \$), the median daily volatility (in percents), the median percentage spread (in basis points), the median monthly turnover rate (in percents). Panel B shows the average and median order size (in percents of daily volume and in basis points of market capitalization) as well as average implementation shortfall (in basis points), the average commission (in basis points), and the average SEC fee for sell orders (in percents per 10 basis points). The thresholds of ten volume groups correspond to 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of dollar volume for common stocks listed on the NYSE. Group 1 (Group 10) contains orders in stocks with lowest (highest) dollar trading volume. The sample ranges from January 2001 to December 2005.

Table 2: OLS Estimates of Order Size.

	NYSE			NASDAQ	
	All	Buy	Sell	Buy	Sell
$\ln [\bar{q}]$	-5.67 (0.017)	-5.68 (0.023)	-5.63 (0.018)	-5.75 (0.035)	-5.65 (0.032)
α_0	-0.62 (0.009)	-0.63 (0.011)	-0.59 (0.008)	-0.71 (0.019)	-0.59 (0.015)
R^2	0.3167	0.2587	0.2646	0.4298	0.3542
$Q^*/V^* \cdot \delta^{-1} \times 10^{-4}$	34.62	34.14	35.98	31.80	34.78
#Obs	439,765	131,530	150,377	69,871	87,987

Table presents the estimates $\ln \bar{q}$ and α_0 for the regression:

$$\ln \left[\frac{X_i}{V_i} \right] = \ln [\bar{q}] + \alpha_0 \cdot \ln \left[\frac{W_i}{W^*} \right] + \tilde{\epsilon}.$$

Each observation corresponds to transition order i with order size X_i , benchmark price $P_{0,i}$, expected daily volume V_i , expected daily volatility σ_i , trading activity W_i . \bar{q} is the measure of order size such that Q^*/V^* measures the median order size for a benchmark stock. The benchmark stock has daily volatility of 2%, share price of \$40, and daily volume of one million shares. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. The sample ranges from January 2001 to December 2005.

Table 3: OLS Estimates for Order Size: Model Calibration.

	All	NYSE		NASDAQ	
		Buy	Sell	Buy	Sell
<i>Restricted Specification: $\alpha_0 = -2/3, b_1 = b_2 = b_3 = b_4 = 0$</i>					
$\ln [\bar{q}]$	-5.71 (0.019)	-5.70 (0.023)	-5.68 (0.019)	-5.70 (0.042)	-5.77 (0.039)
$Q^*/V^* \cdot \delta^{-1} \times 10^{-4}$	33.13	33.46	34.14	33.46	31.20
MSE	2.53	2.61	2.54	2.32	2.56
R^2	0.3149	0.2578	0.2599	0.4278	0.3479
<i>Unrestricted Specification With 5 Degrees of Freedom: $\alpha_0 = -2/3$.</i>					
R^2	0.3229	0.2668	0.2739	0.4318	0.3616
#Obs	439,765	131,530	150,377	69,871	87,987

Table presents the estimates $\ln \bar{q}$ and the mean squared error (MSE) for the regression:

$$\ln \left[\frac{X_i}{V_i} \right] = \ln [\bar{q}] + \alpha_0 \cdot \ln \left[\frac{W_i}{W^*} \right] + b_1 \cdot \ln \left[\frac{\sigma_i}{0.02} \right] + b_2 \cdot \ln \left[\frac{P_{0,i}}{40} \right] + b_3 \cdot \ln \left[\frac{V_i}{10^6} \right] + b_4 \cdot \ln \left[\frac{\nu_i}{1/12} \right] + \tilde{\epsilon}.$$

with α_0 restricted to be $-2/3$ as predicted by invariance hypothesis and $b_1 = b_2 = b_3 = 0$. Each observation corresponds to transition order i with order size X_i , benchmark price $P_{0,i}$, expected daily volume V_i , expected daily volatility σ_i , trading activity W_i , and monthly turnover rate ν_i . \bar{q} is the measure of order size such that Q^*/V^* measures the corresponding percentile of order size for a benchmark stock. The benchmark stock has daily volatility of 2%, share price of \$40, and daily volume of one million shares. The R-squares are reported for restricted specification with $\alpha_0 = -2/3, b_1 = b_2 = b_3 = b_4 = 0$ as well as for unrestricted specification with coefficients $\ln \bar{q}$ and b_1, b_2, b_3, b_4 allowed to vary freely. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. The sample ranges from January 2001 to December 2005.

Table 4: Transactions Cost Estimates in Non-Linear Regression.

	NYSE		NASDAQ		
	All	Buy	Sell	Buy	Sell
β_{mkt}	0.66 (0.013)	0.63 (0.016)	0.62 (0.016)	0.76 (0.037)	0.78 (0.036)
$\bar{\kappa}_0^* \times 10^4$	1.77 (0.837)	-0.27 (2.422)	1.14 (1.245)	0.77 (4.442)	3.55 (1.415)
α_1	-0.49 (0.050)	-0.37 (1.471)	-0.50 (0.114)	0.53 (1.926)	-0.44 (0.045)
$\bar{\kappa}_I^* \times 10^4$	10.69 (1.376)	12.08 (2.693)	9.56 (2.254)	12.33 (2.356)	9.34 (2.686)
z	0.57 (0.039)	0.54 (0.056)	0.56 (0.062)	0.44 (0.051)	0.63 (0.086)
α_2	-0.32 (0.015)	-0.40 (0.037)	-0.33 (0.029)	-0.41 (0.035)	-0.29 (0.037)
R^2	0.1010	0.1118	0.1029	0.0945	0.0919
#Obs	439,765	131,530	150,377	69,871	87,987

Table presents the estimates for β_{mkt} , z , α_1 , $\bar{\kappa}_0^*$, α_2 , and $\bar{\kappa}_I^*$ in the regression:

$$\mathbb{I}_{BS,i} \cdot S_i \cdot \frac{(0.02)}{\sigma_i} = \beta_{mkt} \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} + \mathbb{I}_{BS,i} \cdot \bar{\kappa}_0^* \cdot \left[\frac{W_i}{W^*} \right]^{\alpha_1} + \mathbb{I}_{BS,i} \cdot \bar{\kappa}_I^* \cdot \left[\frac{W_i}{W^*} \right]^{\alpha_2} \cdot \left[\frac{\phi I_i}{0.01} \right]^z + \tilde{\epsilon}. \quad (62)$$

where $\phi I_i / 0.01 = X_i / (0.01 V_i) \cdot (W_i / W^*)^{2/3}$. S_i is implementation shortfall. R_{mkt} is the value-weight market return for the first day of transition. The trading activity W_i is the product of expected volatility σ_i , benchmark price $P_{0,i}$, and expected volume V_i . The scaling constant $W^* = (0.02)(40)(10^6)$ is the trading activity for the benchmark stock with volatility of 2% per day, price \$40 per share, and trading volume of one million shares per day. X_i is the number of shares in the order i . $\bar{\kappa}_I^*$ is the market impact costs of executing a trade of one percent of daily volume in a benchmark stock, and $\bar{\kappa}_0^*$ is the effective spread cost. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. The sample ranges from January 2001 to December 2005.

Table 5: Transactions Costs: Model Calibration.

	All	NYSE		NASDAQ	
		Buy	Sell	Buy	Sell
<i>Linear Model: $z = 1, \beta_1 = \beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = \beta_7 = \beta_8 = 0.$</i>					
$\bar{\kappa}_0^* \times 10^4$	8.2134 (0.5776)	7.1934 (1.1215)	6.7698 (0.7943)	9.1832 (1.5627)	9.2658 (0.7811)
$\bar{\kappa}_I^* \times 10^4$	2.5003 (0.1903)	3.3663 (0.3700)	1.9220 (0.2650)	3.4614 (0.3953)	2.4629 (0.3267)
R^2	0.0991	0.1102	0.1012	0.0926	0.0897
<i>Square Root Model: $z = 1/2, \beta_1 = \beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = \beta_7 = \beta_8 = 0.$</i>					
$\bar{\kappa}_0^* \times 10^4$	2.0763 (0.7035)	-1.3091 (1.2779)	0.9167 (0.9264)	2.2844 (2.0554)	4.6530 (0.8244)
$\bar{\kappa}_I^* \times 10^4$	12.0787 (0.7416)	15.6544 (1.2177)	11.0986 (1.2979)	13.5025 (1.4564)	10.4063 (1.2069)
R^2	0.1007	0.1116	0.1027	0.0941	0.0911
<i>Unrestricted Specification With 12 Degrees of Freedom.</i>					
R^2	0.1016	0.1121	0.1032	0.0957	0.0944
#Obs	439,765	131,530	150,377	69,871	87,987

Table presents the estimates $\bar{\kappa}_0^*$ and $\bar{\kappa}_I^*$ for the regression:

$$\mathbb{I}_{BS,i} \cdot S_i \cdot \frac{(0.02)}{\sigma_i} = \beta_{mkt} \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} \mathbb{I}_{BS,i} \cdot \bar{\kappa}_0^* \cdot \left[\frac{W_i}{W^*} \right]^{-1/3} \cdot \frac{\sigma_i^{\beta_1} \cdot P_{0,i}^{\beta_2} \cdot V_i^{\beta_3} \cdot \nu_i^{\beta_4}}{(0.02)(40)(10^6)(1/12)} + \\ + \mathbb{I}_{BS,i} \cdot \bar{\kappa}_I^* \cdot \left[\frac{\phi I_i}{0.01} \right]^z \cdot \left[\frac{W_i}{W^*} \right]^{-1/3} \cdot \frac{\sigma_i^{\beta_5} \cdot P_{0,i}^{\beta_6} \cdot V_i^{\beta_7} \cdot \nu_i^{\beta_8}}{(0.02)(40)(10^6)(1/12)} + \tilde{\epsilon}.$$

where invariant $\phi I_i/0.01 = X_i/(0.01V_i) \cdot (W_i/W^*)^{2/3}$. S_i is implementation shortfall. R_{mkt} is the value-weight market return for the first day of transition. The trading activity W_i is the product of expected volatility σ_i , benchmark price $P_{0,i}$, and expected volume V_i . The scaling constant $W^* = (0.02)(40)(10^6)$ is the trading activity for the benchmark stock with volatility of 2% per day, price \$40 per share, and trading volume of one million shares per day. X_i is the number of shares in the order i . $\bar{\kappa}_I^*$ is the market impact costs of executing a trade of one percent of daily volume in a benchmark stock, and $\bar{\kappa}_0^*$ is the effective spread cost. The R-squares are reported for restricted specification as well as for unrestricted specification with twelve coefficients $\beta_{mkt}, z, \bar{\lambda}, \bar{\kappa}, \beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8$ allowed to vary freely. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. The sample ranges from January 2001 to December 2005.

Table 6: OLS Estimates of Log of Quoted Spread.

	All	NYSE		NASDAQ	
		Buy	Sell	Buy	Sell
$\ln \bar{s}$	-3.07 (0.008)	-3.09 (0.008)	-3.08 (0.008)	-3.04 (0.013)	-3.04 (0.012)
α_3	-0.35 (0.003)	-0.31 (0.003)	-0.32 (0.003)	-0.40 (0.004)	-0.39 (0.004)
R^2	0.4744	0.3545	0.3964	0.5516	0.5721
$e^{\ln \bar{s}} \cdot 0.02 \times 10^4$	9.28	9.10	9.19	9.57	9.57
#Obs	434,920	130,700	149,197	68,833	86,190

Table presents the estimates $\ln \bar{s}$ and α_3 for the regression:

$$\ln \left[\frac{s_i}{P_i \cdot \sigma_i} \right] = \ln \bar{s} + \alpha_3 \cdot \ln \left[\frac{W_i}{W^*} \right] + \tilde{\epsilon}_i,$$

Each observation corresponds to order i . The left-hand side variable is the logarithm of the quoted bid-ask spread $s_i/P_{0,i}$ as a fraction of expected return volatility σ_i . The trading activity W_i is the product of expected daily volatility σ_i , benchmark price $P_{0,i}$, and expected daily volume V_i , measured as the last month's average daily volume. The scaling constant $W^* = (0.02)(40)(10^6)$ corresponds to the trading activity for the benchmark stock with volatility of 2% per day, price \$40 per share, and trading volume of one million shares per day. The median percentage spread for a benchmark stock is $exp(\ln \bar{s}) \cdot 0.02$. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. The sample ranges from January 2001 to December 2005.

Figure 2: Invariant Order Size Distribution.

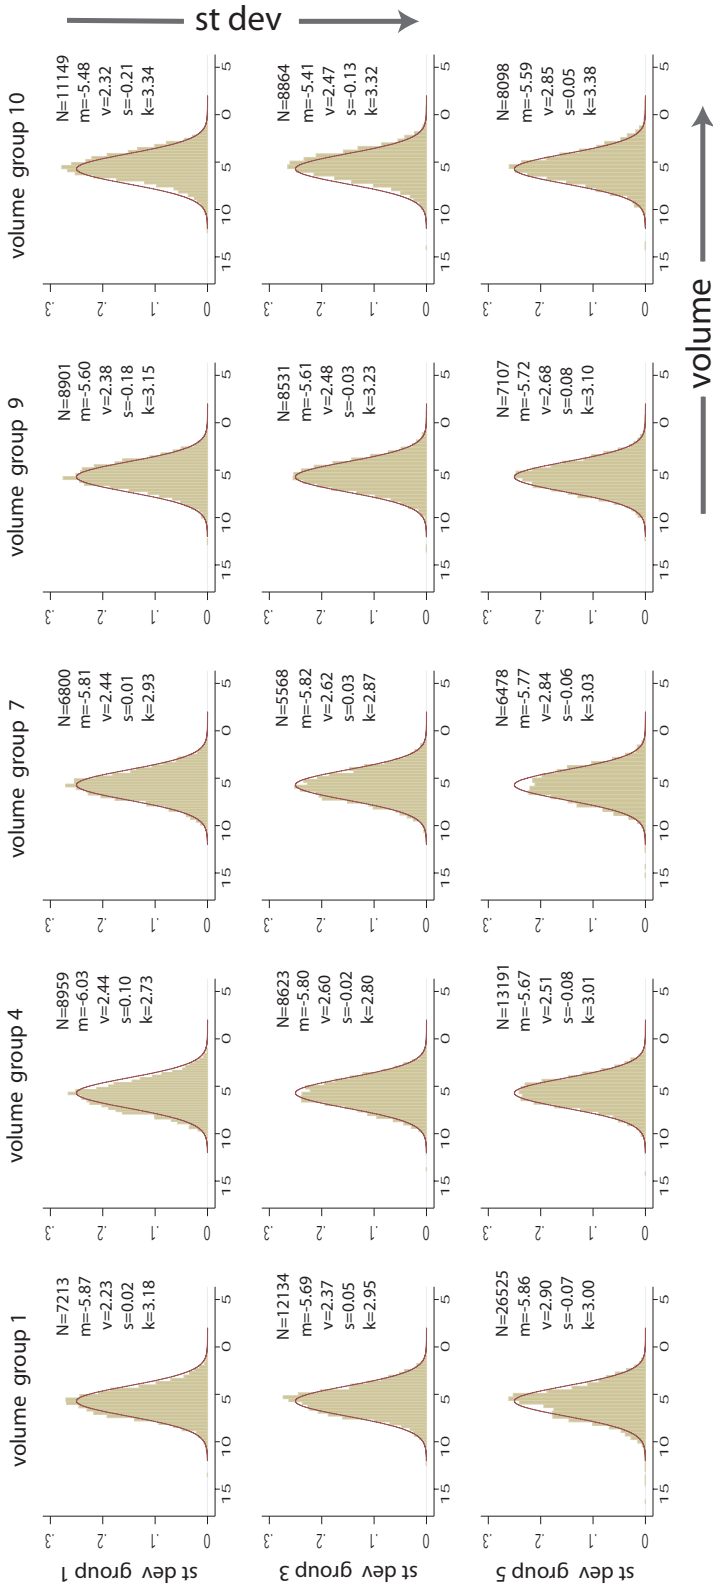




Figure shows distributions of $\ln(\tilde{X}/V) + 2/3 \ln(W_i/W^*)$ for stocks sorted into 10 volume groups and 5 volatility groups (only volume groups 1, 4, 7, 9, 10 and volatility groups 1, 3, 5 are reported). X_i is an order size in shares, V_i is the average daily volume in shares, and W_i is the trading activity. The thresholds of ten volume groups correspond to 30th, 50th, 60th, 70th, 80th, 85th, 90th, and 95th percentiles of dollar volume for common stocks listed on the NYSE. Volume group 1 (group 10) contains orders in stocks with lowest (highest) dollar volume. The thresholds of five volatility groups correspond to 20th, 40th, 60th, and 80th percentiles for common NYSE-listed stocks. Volatility group 1 (group 5) has stocks with the lowest (highest) volatility. Each subplot also shows the number of observations (N), the mean (m), the variance (v), the skewness (s), and the kurtosis (k) for depicted distribution. The normal distribution with the common mean of -5.71 and variance of 2.54 is imposed on each subplot. The common mean and variance are calculated as the mean and variance of distribution over the entire sample. The sample ranges from January 2001 to December 2005.

Figure 3: Invariant Order Size Distribution.

Panel A: Quantile-to-Quantile Plot for Empirical and Lognormal Distribution.



Panel B: Logarithm of Ranks against Quantiles of Empirical Distribution.



Panel A shows quantile-quantile plots of empirical distributions of $\ln(\tilde{X}/V) + 2/3 \ln(W_i/W^*)$ and a normal distribution for stocks sorted into 10 volume groups (only volume groups 1, 4, 7, 9, 10 are reported). Panel B depicts the logarithm of ranks based on that distribution. The ten volume groups are based on average dollar trading volume with thresholds corresponding to 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of the dollar volume for common NYSE-listed stocks. Volume group 1 (group 10) has stocks with the lowest (highest) trading volume. Each subplot shows the number of observations (N), the mean (m), the variance (v), the skewness (s), and the kurtosis (k) of a depicted distribution. There are 400,000+ data points. The sample ranges from January 2001 to December 2005.

Figure 4: Invariant Transactions Cost Functions.

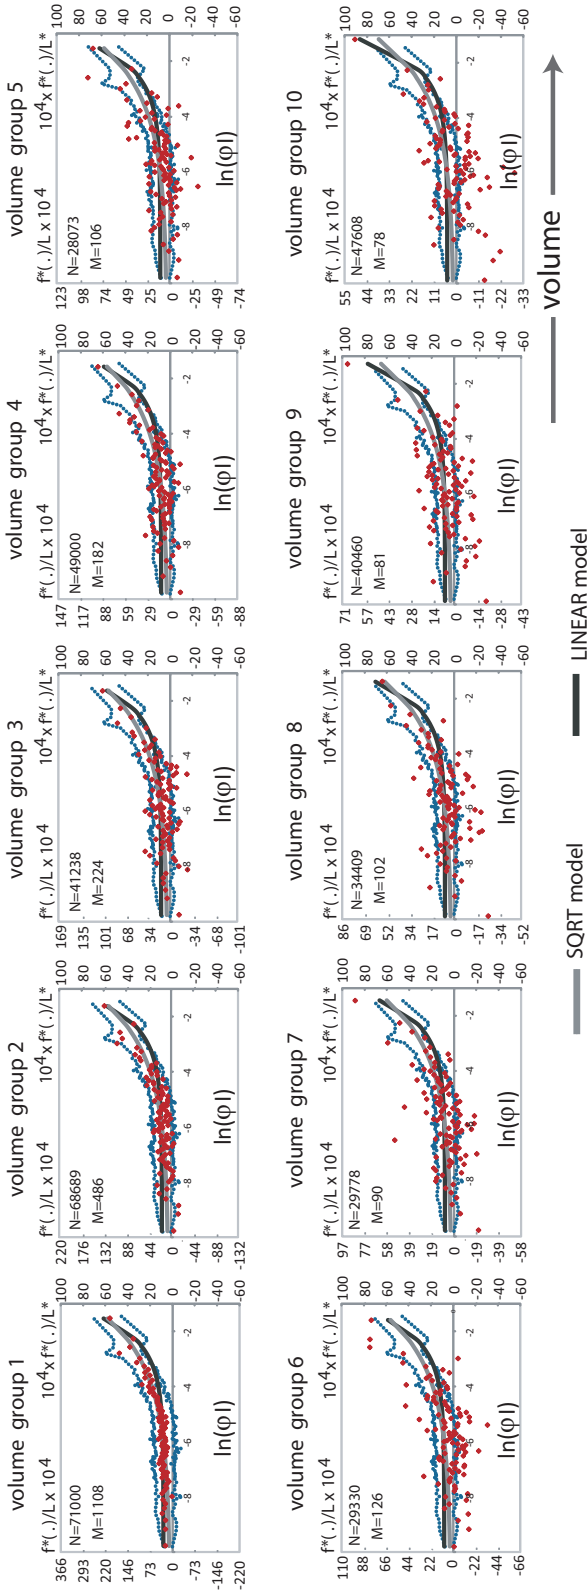


Figure shows estimates of transactions cost functions for stock sorted into 10 volume groups. On the horizontal axis, there are 100 equally spaced bins based on re-scaled order sizes, $\phi I = \tilde{X}/V \cdot (W_i/W^*)^{2/3}$. For each volume group $k = 1, \dots, 10$, subplot contains 100 estimates of dummy variables $f^*(k, h)/L^*$, $h = 1, \dots, 100$ from regression (49). On the right-hand side vertical axis, there are units of scaled transactions cost $f^*(.) / L^*$ for a benchmark stock. On the left-hand side vertical axis, there are units of actual transactions cost $f^*(.) / L^k$ for a benchmark stock, where $1/L^k$ is the illiquidity measure for orders in volume group k . The 95th percent confidence interval estimates based on the entire sample are imposed on each subplot (blue dotted lines). The common linear and square-root functions are imposed on each subplot with the parameter estimated on the entire sample. A linear function is $2.50 \cdot 10^{-4} \cdot \phi I / 0.01 + 8.21 \cdot 10^{-4}$ (black solid line). A square-root function is $12.07 \cdot 10^{-4} \cdot \sqrt{\phi I / 0.01} + 2.08 \cdot 10^{-4}$ (grey solid line). The thresholds of ten volume groups correspond to 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of dollar volume for common stocks listed on the NYSE. Group 1 (group 10) contains orders in stocks with lowest (highest) dollar volume. Each subplot also shows the number of observations N and the number of stocks M (for the last month). The sample ranges from January 2001 to December 2005.

Figure 5: Transactions Cost Functions.




Figure shows estimates of transactions cost functions based on entire sample. On the horizontal axis, there are 100 equally spaced bins based on re-scaled order sizes, $\phi I = \tilde{X}/V \cdot (W_i/W^*)^{2/3}$. The plot contains 100 estimates $f^*(k, h)/L^*$, $h = 1, \dots, 100$ from regression,

$$\mathbb{I}_{BS,i} \cdot S_i \cdot \frac{(0.02)}{\sigma_i} = \beta_{mkt} \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} + \mathbb{I}_{BS,i} \cdot \left[\frac{W_i}{W^*} \right]^{-1/3} \cdot \sum_{h=1}^{100} D_i^*(k, h) \cdot f^*(k, h)/L^* + \tilde{\epsilon}_i.$$

X_i is an order size in shares, V_i is the average daily volume in shares, and W_i is the measure of trading activity. The vertical axis presents estimated transactions cost invariant $f^*(.) / L$ in basis points. The 95th percent confidence interval are superimposed (dotted lines). A linear function is $2.50 \cdot 10^{-4} \cdot \phi I / 0.01 + 8.21 \cdot 10^{-4}$ (black solid line). A square-root function is $12.07 \cdot 10^{-4} \cdot \sqrt{\phi I / 0.01} + 2.08 \cdot 10^{-4}$ (grey solid line). The sample ranges from January 2001 to December 2005.

Part I
Appendix

Table 7: Quantile Estimates of Order Size.

	p1	p5	p25	p50	p75	p95	p99
$\ln [\bar{q}]$	-9.37 (0.008)	-8.31 (0.006)	-6.73 (0.004)	-5.66 (0.003)	-4.59 (0.004)	-3.05 (0.006)	-2.05 (0.009)
α_0	-0.65 (0.005)	-0.64 (0.003)	-0.61 (0.002)	-0.62 (0.002)	-0.61 (0.002)	-0.64 (0.003)	-0.63 (0.005)
Pseudo R^2	0.1621	0.1534	0.1650	0.1727	0.1795	0.1949	0.2232
$Q^*/V^* \cdot \delta \times 10^{-4}$	0.85	2.46	11.95	34.83	101.53	473.59	1287.35
#Obs	439,765	439,765	439,765	439,765	439,765	439,765	439,765

Table presents the estimates $\ln \bar{q}$ and α_0 for the quantile regression:

$$\ln \left[\frac{X_i}{V_i} \right] = \ln [\bar{q}] + \alpha_0 \cdot \ln \left[\frac{W_i}{W^*} \right] + \tilde{\epsilon}.$$

Each observation corresponds to transition order i with order size X_i , benchmark price $P_{0,i}$, expected daily volume V_i , expected daily volatility σ_i , trading activity W_i . \bar{q} is the measure of order size such that Q^*/V^* measures the corresponding percentile of order size for a benchmark stock. The benchmark stock has daily volatility of 2%, share price of \$40, and daily volume of one million shares. The standard errors are shown in parentheses. The sample ranges from January 2001 to December 2005.

Table 8: OLS Estimates for Order Size: Model Calibration.

	All	NYSE		NASDAQ	
		Buy	Sell	Buy	Sell
<i>Restricted Specification: $\alpha_0 = -2/3, b_1 = b_2 = b_3 = b_4 = 0$</i>					
$\ln [\bar{q}]$	-5.71 (0.019)	-5.70 (0.023)	-5.68 (0.019)	-5.70 (0.042)	-5.77 (0.039)
$Q^*/V^* \cdot \delta \times 10^4$	33.13	33.46	34.14	33.46	31.20
MSE	2.53	2.61	2.54	2.32	2.56
R^2	0.3149	0.2578	0.2599	0.4278	0.3479
<i>Unrestricted Specification With 5 Degrees of Freedom: $\alpha_0 = -2/3$.</i>					
$\ln [\bar{q}]$	-5.53 (0.019)	-5.55 (0.026)	-5.48 (0.019)	-5.77 (0.051)	-5.48 (0.047)
b_1	0.42 (0.040)	0.47 (0.050)	0.53 (0.043)	0.19 (0.094)	0.33 (0.087)
b_2	0.24 (0.019)	0.17 (0.021)	0.29 (0.017)	0.04 (0.049)	0.33 (0.040)
b_3	0.06 (0.010)	0.06 (0.012)	0.07 (0.009)	-0.06 (0.026)	0.07 (0.021)
b_4	-0.18 (0.015)	-0.24 (0.020)	-0.22 (0.017)	-0.02 (0.040)	-0.11 (0.032)
R^2	0.3229	0.2668	0.2739	0.4318	0.3616
#Obs	439,765	131,530	150,377	69,871	87,987

Table presents the estimates and the mean squared error (MSE) for the regression:

$$\ln \left[\frac{X_i}{V_i} \right] = \ln [\bar{q}] + \alpha_0 \cdot \ln \left[\frac{W_i}{W^*} \right] + b_1 \cdot \ln \left[\frac{\sigma_i}{0.02} \right] + b_2 \cdot \ln \left[\frac{P_{0,i}}{40} \right] + b_3 \cdot \ln \left[\frac{V_i}{10^6} \right] + b_4 \cdot \ln \left[\frac{\nu_i}{1/12} \right] + \tilde{\epsilon}.$$

with α_0 restricted to be $-2/3$ as predicted by invariance hypothesis and $b_1 = b_2 = b_3 = 0$. Each observation corresponds to transition order i with order size X_i , benchmark price $P_{0,i}$, expected daily volume V_i , expected daily volatility σ_i , trading activity W_i , and monthly turnover rate ν_i . \bar{q} is the measure of order size such that Q^*/V^* measures the corresponding percentile of order size for a benchmark stock. The benchmark stock has daily volatility of 2%, share price of \$40, and daily volume of one million shares. The R-squares are reported for restricted specification with $\alpha_0 = -2/3, b_1 = b_2 = b_3 = b_4 = 0$ as well as for unrestricted specification with coefficients $\ln \bar{q}$ and b_1, b_2, b_3, b_4 allowed to vary freely. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. The sample ranges from January 2001 to December 2005.

Table 9: Transactions Cost Estimates in Non-Linear Regression with Linear Impact.

	NYSE			NASDAQ	
	All	Buy	Sell	Buy	Sell
β_{mkt}	0.66 (0.013)	0.63 (0.016)	0.62 (0.016)	0.77 (0.037)	0.77 (0.036)
$\bar{\kappa}_0^* \times 10^4$	6.28 (0.890)	6.51 (1.600)	5.43 (1.154)	5.94 (2.147)	6.54 (1.501)
α_1	-0.40 (0.020)	-0.36 (0.048)	-0.39 (0.029)	-0.44 (0.051)	-0.40 (0.031)
$\bar{\kappa}_I^* \times 10^4$	2.73 (0.252)	2.63 (0.460)	2.10 (0.346)	3.69 (0.663)	3.13 (0.765)
α_2	-0.31 (0.028)	-0.45 (0.038)	-0.31 (0.041)	-0.32 (0.056)	-0.28 (0.058)
R^2	0.0993	0.1105	0.1014	0.0931	0.0901
#Obs	439,765	131,530	150,377	69,871	87,987

Table presents the estimates for β_{mkt} , α_1 , $\bar{\kappa}_0^*$, α_2 , and $\bar{\kappa}_I^*$ in the regression:

$$\mathbb{I}_{BS,i} \cdot S_i \cdot \frac{(0.02)}{\sigma_i} = \beta_{mkt} \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} + \mathbb{I}_{BS,i} \cdot \bar{\kappa}_0^* \cdot \left[\frac{W_i}{W^*} \right]^{\alpha_1} + \mathbb{I}_{BS,i} \cdot \bar{\kappa}_I^* \cdot \left[\frac{W_i}{W^*} \right]^{\alpha_2} \cdot \left[\frac{\phi I_i}{0.01} \right]^z + \tilde{\epsilon}. \quad (63)$$

where $z = 1$ and $\phi I_i / 0.01 = X_i / (0.01 V_i) \cdot (W_i / W^*)^{2/3}$. S_i is implementation shortfall. R_{mkt} is the value-weight market return for the first day of transition. The trading activity W_i is the product of expected volatility σ_i , benchmark price $P_{0,i}$, and expected volume V_i . The scaling constant $W^* = (0.02)(40)(10^6)$ is the trading activity for the benchmark stock with volatility of 2% per day, price \$40 per share, and trading volume of one million shares per day. X_i is the number of shares in the order i . $\bar{\kappa}_I^*$ is the market impact costs of executing a trade of one percent of daily volume in a benchmark stock, and $\bar{\kappa}_0^*$ is the effective spread cost. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. The sample ranges from January 2001 to December 2005.

Table 10: Transactions Cost Estimates in Non-Linear Regression with Quoted Spread.

	All	NYSE		NASDAQ	
		Buy	Sell	Buy	Sell
β_{mkt}	0.65 (0.013)	0.63 (0.016)	0.62 (0.015)	0.76 (0.036)	0.77 (0.037)
$\bar{\kappa}_I \times 10^4$	2.95 (0.261)	2.97 (0.504)	2.24 (0.366)	3.76 (0.700)	2.95 (0.749)
α_2	-0.32 (0.029)	-0.44 (0.036)	-0.32 (0.039)	-0.37 (0.053)	-0.33 (0.060)
h	0.71 (0.053)	0.61 (0.110)	0.74 (0.094)	0.61 (0.127)	0.75 (0.073)
R^2	0.0976	0.1094	0.1010	0.0891	0.0872
#Obs	436,649	131,100	149,600	69,218	86,731

Table presents the estimates for β_{mkt} , $\bar{\kappa}_I$, α_2 , and h in the regression:

$$\mathbb{I}_{BS,i} \cdot S_i \cdot \frac{(0.02)}{\sigma_i} = \beta_{mkt} \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} + \mathbb{I}_{BS,i} \cdot h \cdot \frac{1}{2} \cdot \frac{s_i}{P_{0,i}} \cdot \frac{(0.02)}{\sigma_i} + \mathbb{I}_{BS,i} \cdot \bar{\kappa}_I \cdot \left[\frac{\phi I_i}{0.01} \right] \cdot \left[\frac{W_i}{W^*} \right]^{\alpha_2} + \tilde{\epsilon}.$$

where invariant $I_i = \frac{X_i}{(0.01)V_i} \cdot \left[\frac{W_i}{W^*} \right]^{2/3}$. Each observation corresponds to order i . $\mathbb{I}_{BS,i}$ is a buy/sell indicator, S_i is implementation shortfall, R_{mkt} is the value-weight market return for the first day of transition. The term $(0.02)/\sigma_i$ adjusts for heteroscedasticity. The trading activity W_i is the product of expected volatility σ_i , benchmark price $P_{0,i}$, and expected volume V_i . The scaling constant $W^* = (0.02)(40)(10^6)$ is the trading activity for the benchmark stock with volatility of 2% per day, price \$40 per share, and trading volume of one million shares per day. X_i is the number of shares in the order i . $\bar{\kappa}_I$ is the market impact costs of executing a trade of one percent of daily volume in a benchmark stock. $s_i/P_{0,i}$ is the quoted spread. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. The sample ranges from January 2001 to December 2005.

Table 11: Transactions Costs: Model Calibration.

	NYSE		NASDAQ		
	All	Buy	Sell	Buy	Sell
<i>Linear Model: $z = 1, \beta_1 = \beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = \beta_7 = \beta_8 = 0.$</i>					
β_{mkt}	0.6571 (0.0135)	0.6308 (0.0159)	0.6195 (0.0158)	0.7693 (0.0371)	0.7771 (0.0365)
$\bar{\kappa}_0 \times 10^4$	8.2134 (0.5776)	7.1934 (1.1215)	6.7698 (0.7943)	9.1832 (1.5627)	9.2658 (0.7811)
$\bar{\kappa}_I \times 10^4$	2.5003 (0.1903)	3.3663 (0.3700)	1.9220 (0.2650)	3.4614 (0.3953)	2.4629 (0.3267)
R^2	0.0991	0.1102	0.1012	0.0926	0.0897
<i>Square Root Model: $z = 1/2, \beta_1 = \beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = \beta_7 = \beta_8 = 0.$</i>					
β_{mkt}	0.6552 (0.0134)	0.6285 (0.0158)	0.6192 (0.0159)	0.7598 (0.0365)	0.7782 (0.0364)
$\bar{\kappa}_0 \times 10^4$	2.0763 (0.7035)	-1.3091 (1.2779)	0.9167 (0.9264)	2.2844 (2.0554)	4.6530 (0.8244)
$\bar{\kappa}_I \times 10^4$	12.0787 (0.7416)	15.6544 (1.2177)	11.0986 (1.2979)	13.5025 (1.4564)	10.4063 (1.2069)
R^2	0.1007	0.1116	0.1027	0.0941	0.0911
<i>Unrestricted Specification With 12 Degrees of Freedom.</i>					
β_{mkt}	0.66 (0.013)	0.63 (0.016)	0.62 (0.015)	0.76 (0.036)	0.78 (0.036)
$\bar{\kappa}_0^* \times 10^4$	0.94 (0.675)	-0.05 (0.124)	0.47 (0.556)	1.55 (1.698)	1.61 (1.148)
β_1	-0.43 (0.147)	-2.47 (0.890)	-1.08 (0.392)	-0.44 (0.489)	-0.46 (0.131)
β_2	0.17 (0.072)	2.87 (1.230)	0.23 (0.231)	0.20 (0.127)	0.11 (0.109)
β_3	-0.56 (0.159)	1.85 (0.754)	-0.47 (0.296)	-0.47 (0.238)	-0.49 (0.155)
β_4	0.62 (0.173)	0.13 (0.620)	0.49 (0.490)	0.49 (0.313)	0.58 (0.175)
$\bar{\kappa}_I^* \times 10^4$	9.36 (1.307)	11.61 (2.471)	10.93 (1.804)	8.88 (3.340)	5.00 (2.033)
z	0.58 (0.041)	0.54 (0.039)	0.52 (0.042)	0.58 (0.094)	0.63 (0.083)
β_5	0.02 (0.135)	-0.11 (0.192)	0.36 (0.229)	-0.17 (0.252)	-0.23 (0.242)
β_6	-0.14 (0.061)	-0.11 (0.113)	0.03 (0.100)	-0.27* (0.120)	-0.22 (0.113)
β_7	0.01 (0.037)	-0.07 (0.050)	0.04 (0.052)	0.00 (0.099)	-0.16 (0.100)
β_8	0.08 (0.067)	0.07 (0.086)	-0.11 (0.101)	0.08 (0.143)	0.39 (0.153)
R^2	0.1016	0.1121	0.1032	0.0957	0.0944
#Obs	439,765	131,530	150,377	69,871	87,987

Table presents the estimates for the regression:

$$\begin{aligned} \mathbb{I}_{BS,i} \cdot S_i \cdot \frac{(0.02)}{\sigma_i} = & \beta_{mkt} \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} \mathbb{I}_{BS,i} \cdot \bar{\kappa}_0^* \cdot \left[\frac{W_i}{W^*} \right]^{-1/3} \cdot \frac{\sigma_i^{\beta_1} \cdot P_{0,i}^{\beta_2} \cdot V_i^{\beta_3} \cdot \nu_i^{\beta_4}}{(0.02)(40)(10^6)(1/12)} + \\ & + \mathbb{I}_{BS,i} \cdot \bar{\kappa}_I^* \cdot \left[\frac{\phi I_i}{0.01} \right]^z \cdot \left[\frac{W_i}{W^*} \right]^{-1/3} \cdot \frac{\sigma_i^{\beta_5} \cdot P_{0,i}^{\beta_6} \cdot V_i^{\beta_7} \cdot \nu_i^{\beta_8}}{(0.02)(40)(10^6)(1/12)} + \tilde{\epsilon}. \end{aligned}$$

where $\phi I_i/0.01 = X_i/(0.01V_i) \cdot (W_i/W^*)^{2/3}$. S_i is implementation shortfall. R_{mkt} is the value-weight market return for the first day of transition. The trading activity W_i is the product of expected volatility σ_i , benchmark price $P_{0,i}$, and expected volume V_i . The scaling constant $W^* = (0.02)(40)(10^6)$ is the trading activity for the benchmark stock with volatility of 2% per day, price \$40 per share, and trading volume of one million shares per day. X_i is the number of shares in the order i . $\bar{\kappa}_I^*$ is the market impact costs of executing a trade of one percent of daily volume in a benchmark stock, and $\bar{\kappa}_0^*$ is the effective spread cost. The R-squares are reported for restricted specification as well as for unrestricted specification with twelve coefficients β_{mkt} , z , $\bar{\kappa}_I^*$, $\bar{\kappa}_0^*$, β_1 , β_2 , β_3 , β_4 , β_5 , β_6 , β_7 , β_8 allowed to vary freely. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. The sample ranges from January 2001 to December 2005.